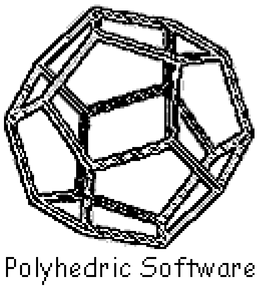

Welcome to Acid WAV's Help system!

At a glance

Acid WAV is an advanced sound editor and synthesizer created by Tommy Anderberg and brought to you by Polyhedric Software. Unlike older sound editors, it has been designed from the ground up to take full advantage of Microsoft's 32 bit Windows architecture, and features

- an unparalleled set of powerful synthesis and editing functions,
- an intuitive, friendly user interface; and
- an incredible price!

If you are running Acid WAV for the first time, please read the License Agreement, then get to know Acid WAV!

Operations

- Loading and saving
- Playing and recording
- Analyzing sound data
- Basic editing
- Advanced editing
- Synthesizing new sounds
- Setting program options
- Automating tasks with scripts

Ordering Acid WAV

- Ordering overview
- Order online

- Order offline

Getting tech support

Submit error reports and questions to support@polyhedric.com.

Before sending us a question, please make sure that it isn't answered in this document. Don't forget to use the Find button - it's much faster than an e-mail query. Our own response time depends mainly on the number of redundant questions we have to handle.

Polyhedric Software

http://www.polyhedric.com

Tommy Anderberg is the creator of several popular audio tools for MS Windows (WAVmaker, MIDInight Express, Mellosoftron, Virtual Sampler SDK). A physicist by training and a computer geek by vocation, he is also a consultant and a prolific freelance writer on personal computing and the Internet.

Tommy can be reached at

Tommy.Anderberg@polyhedric.com

Visit Tommy's homepage at http://www.abc.se/~m9303!

License Agreement

Last updated: February 20, 1998.

This section contains important license information regarding the distribution and use of Acid WAV.

 This software and all the accompanying data and documentation files are protected by national and European laws and by International Treaty provisions. You must treat Acid WAV like any other copyrighted material, except that you may copy the Acid WAV Evaluation Package and give to others (subject to the limitations outlined below). Any use or distribution of this software, data and documentation in violation of Copyright law or of the terms outlined in this document will be prosecuted to the best of our ability.

The conditions under which you may copy and distribute Acid WAV are described below under Limited Distribution License Agreement.

The conditions under which you may use Acid WAV are described below under Limited User License Agreement.

All rights not expressly granted here are reserved to Tommy Anderberg.

Limited Distribution License Agreement

This Limited Distribution License Agreement (throughout this section, the "Agreement") is a legal agreement between you, the distributor, and Tommy Anderberg. By copying any part of Acid WAV, you are agreeing to be bound by the terms of this Agreement.

As the copyright and trademark holder for Acid WAV, Tommy Anderberg authorizes individuals and organizations to copy the files contained in Acid WAV Evaluation Package for their own use or for giving to others, ONLY when the following conditions are met:

- The Acid WAV Evaluation Package - including all related program, data and documentation files - may not be modified in any way and must be distributed together, as a complete package, without exception. You may not distribute corrupted or otherwise damaged files.

- No price or other compensation may be charged for the Acid WAV Evaluation Package. A reasonable distribution cost may be charged for the distribution media, shipping and handling. BBS systems may not charge beyond their normal rates for downloads of Acid WAV Evaluation Package.

- Acid WAV is designed to be easily extended by the addition of patches and script files. You may not, however, distribute an extended version of Acid WAV Evaluation Package without prior written permission from Tommy Anderberg.

- Neither this Agreement nor any part or portion hereof shall be assigned or sublicensed, except as described herein. Should any provision of this Agreement be held to be void, invalid, unenforceable or illegal by a court, the validity and enforceability of the other provisions shall not be affected thereby. If any provision is determined to be unenforceable, you agree to a modification of such provision to provide for enforcement of the provision's intent, to the extent permitted by applicable law. Failure of a party to enforce any provision of this Agreement shall not constitute or be construed as a waiver of such provision or of the right to enforce such provision. If you fail to comply with any terms of this Agreement, your license is automatically terminated, and you may be prosecuted at any time under national and European laws and under International Treaty provisions.

- You acknowledge that you have read this Agreement, you understand this agreement, and understand that by copying any part of Acid WAV, you agree to be bound by this Agreement's terms and conditions. You further agree that, except for written separate agreements between Tommy Anderberg and you, this Agreement is a complete and exclusive statement of your distribution rights. This agreement supersedes all prior oral agreements, proposals or understandings, and any other communications between Tommy Anderberg, Polyhedric Software and you relating to the subject matter of this Agreement.

Limited User License Agreement

This Limited User License Agreement (throughout this section, the "Agreement") is a legal agreement between you, the user, and Tommy Anderberg. By installing Acid WAV, by loading or running any part of it, or by placing or copying any part of it onto your computer hard drive, you are agreeing to be bound by the terms of this Agreement.

For the purposes of this section, "use" means loading the software into RAM, as well as installation on a hard disk or other storage device.

As the copyright and trademark holder for Acid WAV, Tommy Anderberg authorizes individuals or organizations to use Acid WAV, ONLY when the following conditions are met:

- You shall not: modify, translate, disassemble, decompile or reverse engineer the compiled executables contained in Acid WAV.

- You may not exploit running copies of Acid WAV commercially, e.g. by renting, leasing or selling access to a computer running Acid WAV.

- You may not distribute or otherwise exploit the output produced by Acid WAV if you have not purchased your own personal Acid WAV key from Polyhedric Software.

- Acid WAV keys are personal and non-transferable. They may not be given, sold, or otherwise made available to third parties, e.g. (but not limited to) by uploading them to BBS systems or Internet sites.

- The user accepts sole and full responsibility for any damage or loss arising from the use or the inability to use Acid WAV. While this software has been extensively tested and is believed to be safe, Tommy Anderberg and Polyhedric Sowftare disclaim any warranties, either expressed or implied, including but not limited to warranties of fitness for ANY particular purpose and of uninterrupted, error free operation. Under no circumstances shall Tommy Anderberg or Polyhedric Software be liable for loss of data, loss of profits, lost savings, special, incidental, consequential, indirect or other similar damages arising from the use or the inability to use Acid WAV.

- Neither this Agreement nor any part or portion hereof shall be assigned or sublicensed, except as described herein. Should any provision of this Agreement be held to be void, invalid, unenforceable or illegal by a court, the validity and enforceability of the other provisions shall not be affected thereby. If any provision is determined to be unenforceable, you agree to a modification of such provision to provide for enforcement of the provision's intent, to the extent permitted by applicable law. Failure of a party to enforce any provision of this Agreement shall not constitute or be construed as a waiver of such provision or of the right to enforce such provision. If you fail to comply with any terms of this Agreement, your license is automatically terminated, and you may be prosecuted at any time under national and European laws and under International Treaty provisions.

- You acknowledge that you have read this Agreement, you understand this Agreement, and understand that by installing any part of Acid WAV, by loading or running any part of Acid WAV, or by placing or copying any part of Acid WAV onto your computer hard drive, you agree to be bound by this Agreement's terms and conditions. You further agree that, except for written separate agreements between Tommy Anderberg and you, this Agreement is a complete and exclusive statement of the rights and liabilities of the parties involved in the use of Acid WAV. This agreement supersedes all prior oral agreements, proposals or understandings, and any other communications between Tommy Anderberg, Polyhedric Software and you relating to the subject matter of this Agreement.

Getting to know Acid WAV

[Click Contents for all topics.]

When using Acid WAV, you will spend most of your time looking at the main window.

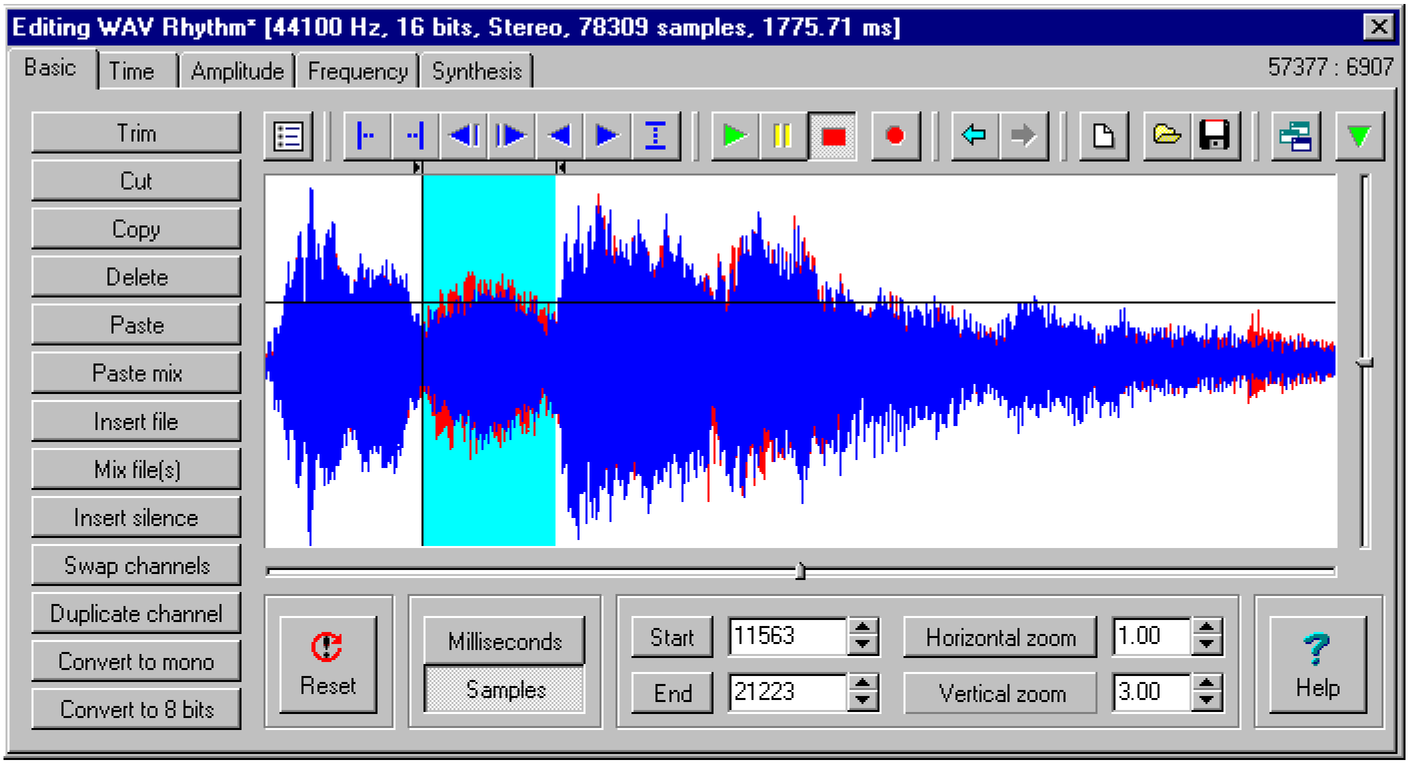

A surrealistic encounter between this window and Jack the Ripper would leave the following objects scattered on your screen:

Title bar

Function tabs

Toolbar
Position display
Section markers
Wave display
Trackbars

Function buttons
Reset button
Time unit box
Selection box
Help button

Their uses include...

Basic editing
Automating tasks
Playing and recording
Loading and saving

Advanced editing
Analyzing sound data

as well as

Synthesizing new sounds

This title bar tells us that we are editing a WAV file called "Rhythm". The asterisk trailing the filename means that the file has been modified since it was last saved. The basic properties of the file are also displayed [between square brackets; in this picture we are editing a CD quality file].

Immediately below the title bar are five tabs and (at the far right, below the window's  button) a cursor position display.

Each tab corresponds to a program function group. In older sound editors, the tabs would be items in a menu bar, and the actual functions would be displayed in drop-down menus brought up by clicking the menu bar. In Acid WAV, the functions in the selected group are always visible and readily available to the left of the graph.

The cursor position display below the main window's  button tells us the mouse location in the graph. The numbers are read as "time position : sample value". The time unit can be either milliseconds or samples, and is selected in the time unit box below the graph.

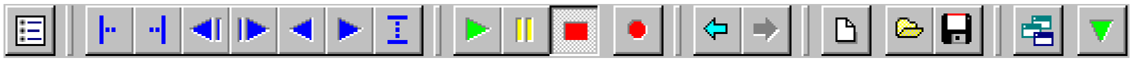

Below the function tabs and the position display is a
toolbar subdivided into six different button groups:

 Script group. Use this button to run scripts.

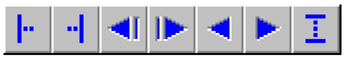 Graph group. Use these buttons to scroll the graph over the loaded file.

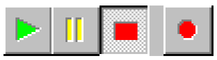 Play/record group. Use these buttons to play and record WAV files.

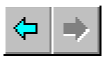 History group. Use these buttons to undo/redo operations.

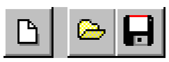 File group. Use these buttons to clear, load and save files.

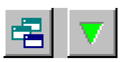 Window group. Use these buttons to open and minimize Acid WAV windows.

 This is the smallest speedbutton "group". Left-clicking the script button causes the default script to be executed on the selected file section. If there is no default script, the Scripts window is opened. Right-clicking the script button also brings up the **Scripts** window.

Scripting is a relatively advanced subject better left alone until you have mastered the rest of Acid WAV.

Use these buttons to scroll
the graph over the loaded
WAV file.

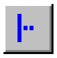 Jump to start of file.

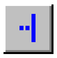 Jump to end of file.

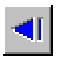 Step back by 10% of file size.

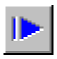 Step forward by 10% of file size.

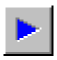 Step forward by one frame.

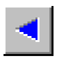 Step back by one frame.

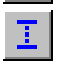 Center horizontal axis.

Use these buttons to play
and record WAV files.

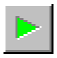 Play the selected file section.

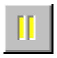 Pause / resume playback.

 Stop playback or recording.

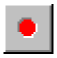 Record. Right-click to change record settings.

Playing and recording is
covered in detail here.

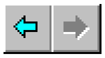 These buttons let you move back and forth between successive versions of the file being edited, just like the back/forward buttons on your web browser let you move between web pages you have visited.

Right-click either button (even if it's greyed out) to bring up the temporary storage window. If the button is enabled, right-clicking it also presents you with the option to clear its "memory" (purge undos and redos, respectively).

Use these buttons to clear the editor
and to load and save WAV files.

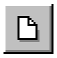 Clear (unload the current file).

 Load file. Right-click to change open settings.

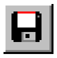 Save file. Right-click to change output format.

Use these buttons to open and minimize Acid WAV windows.

 Open a new Acid WAV window.

Working with multiple Acid WAV windows allows you to assign different default settings to the same functions, edit files independently and then combine the results using the clipboard (see the section on basic editing functions).

Note that each window has its own button on the Windows task bar, labeled with the name of the file being edited.

 Minimize window.

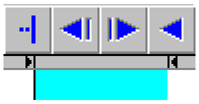

The markers below the speedbuttons are visible only when a section of the loaded file has been selected. They can dragged with the mouse to resize the selected section. Dragging the start marker to a position more than one sample beyond the end marker, or the end marker to a position preceding the start marker by more than one sample, causes the section to be unselected.

Like the graph group buttons, the horizontal trackbar below the wave display is used to scroll through the loaded file.

The vertical trackbar (to the right of the wave display) also allows you to scroll vertically, i.e. to offset the horizontal axis (silence level) in the graph.

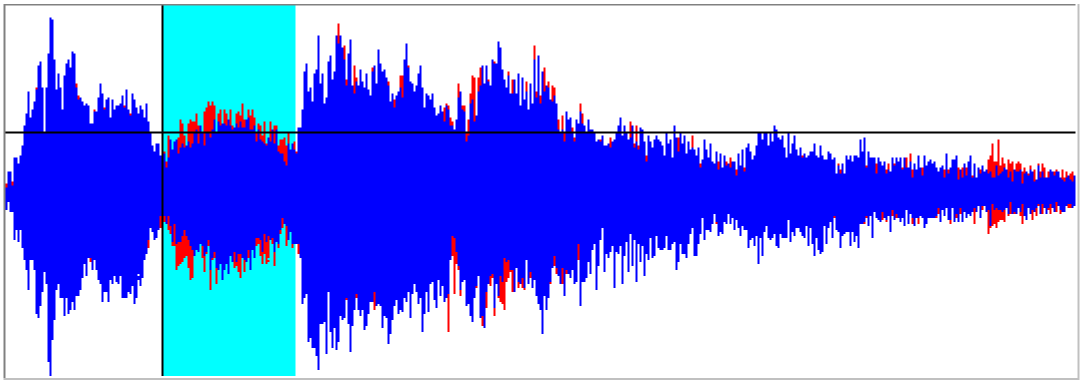

The wave display is politically correct: the left channel is red and the right channel is blue. Monophonic files are displayed in red. If a file section has been selected it's displayed on an aqua-colored background.

Right-clicking the graph brings up the analysis pop-up menu

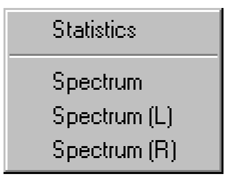

from which you can bring up statistics and frequency spectra for the selected section.

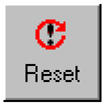

Clicking the **Reset** button causes the wave display to be restored to its default state, as if the file had just been loaded.

- Any selection is cleared.
- The view is scrolled back to the first sample.
- The frame length is set to the default value (see below) or to the file length, whichever is shorter.
- The vertical zoom in the selection box is reset to 1.0 and the horizontal axis in the graph is centered, casing the full amplitude range (0 to 255 for 8 bit files, -32768 to 32767 for 16 bit files) to be displayed.

When the graph is being redrawn, the **Reset** button becomes a **Stop** button, allowing you to abort lengthy redraws.

Right-click this button to set the default open / reset frame length.

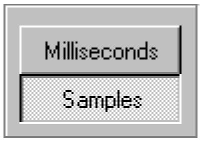

The time unit box located below the main window's wave display is where you set the current time unit.

The selection box is used to...

- Set and modify the **Start** and **End** positions of the selected file section, as described in detail here.

- Center the wave display on the **Start** or **End** positions of the selected file section (just click the corresponding buttons).

- Modify the horizontal zoom (i.e. the time scale) of the wave display. A horizontal zoom of 1.00 means that the whole file is visible; a horizontal zoom of 100.00 means that the graph shows 1% of the file. Clicking the **Horizontal zoom** button causes the display to be fitted exactly to the selected section.

- Modify the vertical zoom (i.e. the amplitude scale) of the wave display. A vertical zoom of 1.00 means that the full amplitude range is visible (0 to 255 for 8 bit files, -32768 to 32767 for 16 bit files).

Clicking the **Help** button opens this document. Right-clicking it brings up the About dialog with information about free memory, system resources and available disk space:

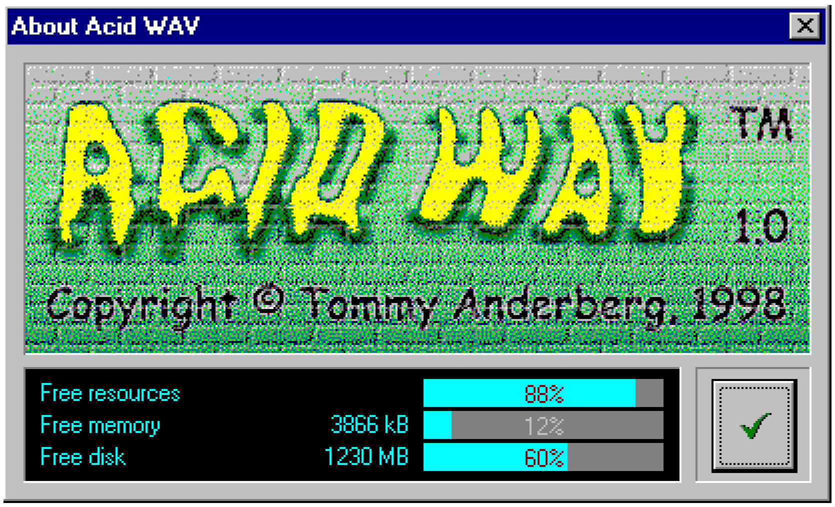

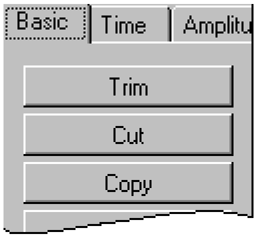

This is where the real action is!

Clicking a function button causes immediate execution of the function on the selected file section (functions which by their nature can only be applied to whole files, e.g. 8 to 16 bit conversion, simply ignore selections).

Right-clicking a function button calls up the function's settings window (provided that the function has one - most do).

Acid WAV remembers all function settings, even between sessions. This means that unlike older sound editors, it only requires you to deal with function settings when you want to change them. If you are happy with your current settings, all you have to do is click!

When no file is loaded, most function buttons will do nothing if clicked (although you can still right-click them to edit function settings). The exceptions are marked with a red triangle:

Apart from Paste, Insert silence and all Synthesis buttons will work even if no file is loaded.

Selecting file sections

See also: standard insertion rules.

No selection

When neither start nor end positions are marked in the wave display, operations are executed on the whole file. In other words, "no selection" implicitly means "all selected".

In this state, the graph cursor is a small cross-hair:

Below the graph, the **End** button is grayed out, and the (horizontal) cursor position is displayed in the **Start** edit box:

Clicking the **Start** button causes the view to be centered on the position in the **Start** edit box.

No end position

Clicking the wave display when neither start nor end positions are marked causes a start marker and a cross-hair to be inserted:

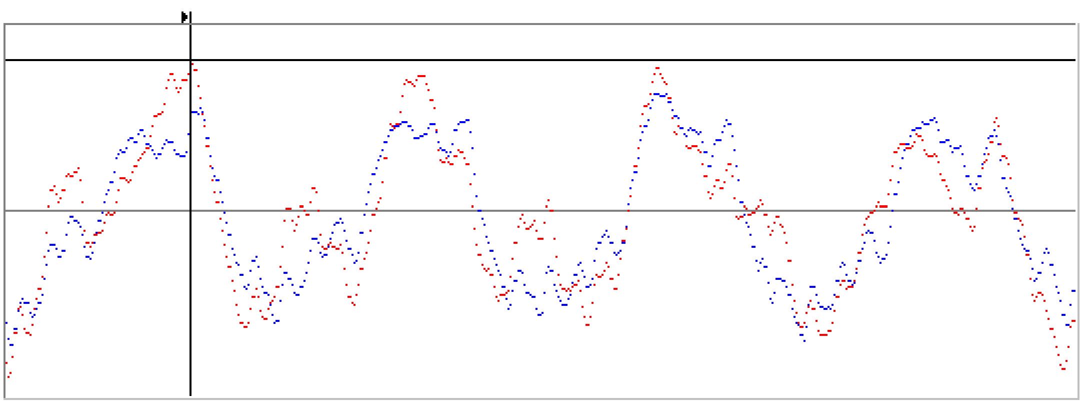

When only a start position is marked in the wave display, operations are executed from that position to the end of the file. In other words, "no end position" implicitly means "to end of file".

In this state, the graph cursor is a horizontal, bidirectional arrow (a "resize" cursor):

Below the graph, the **End** button is activated, and the (horizontal) cursor position is displayed in the **End** edit box. The contents of the **Start** edit box are frozen, reflecting the cross-hair's horizontal position:

You can reposition the cross-hair by editing the value in the **Start** edit box, and center the wave display on it by clicking the **Start** button. The cross-hair can also be repositioned by dragging the start marker above the graph.

Clicking the graph at a position preceding the start marker by more than one sample causes the start marker and the cross-hair to be deleted.

Full selection

Clicking the wave display when only a start position is marked causes an end marker to appear at the clicked position and the section between the two markers to be highlighted:

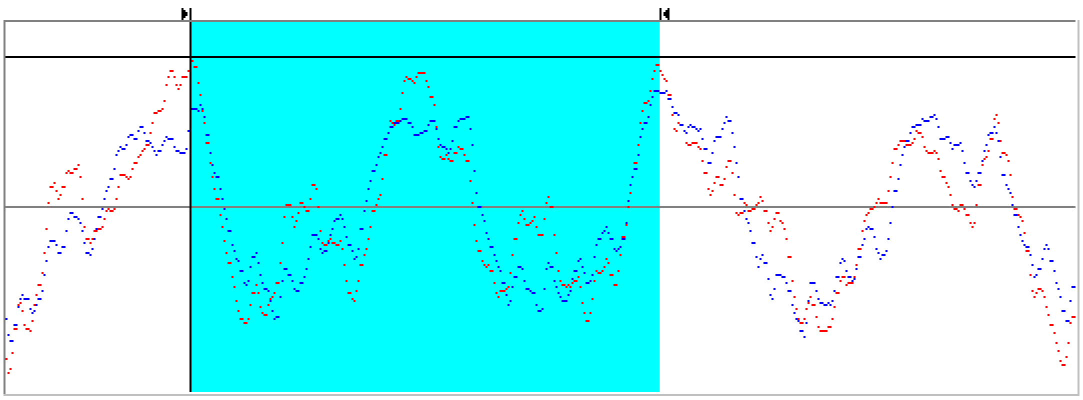

In this state, the graph cursor is a standard Windows arrow: 

Below the graph, the contents of the **Start** and **End** edit boxes are frozen, reflecting the start and end marker's positions. You can move and resize your selection by editing the values in these boxes. The display can be centered on either position by clicking the corresponding button. You can also resize and move the selected section by dragging the start and end markers above the graph. The **Horizontal zoom** button is enabled. Clicking it causes the wave display to be fitted exactly to the selected section.

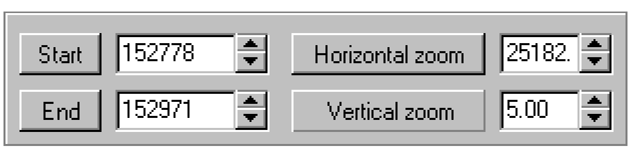

Clicking the graph anywhere, moving the start position more than one sample beyond the end position or moving the end position more than one sample behind the start position causes the marked section to be unselected. Depending on the horizontal zoom settings, clicking **Horizontal zoom** may have the same effect.

Loading and saving

Loading

In order to load a file, click the button in the toolbar. This brings up a standard Windows Open dialog:

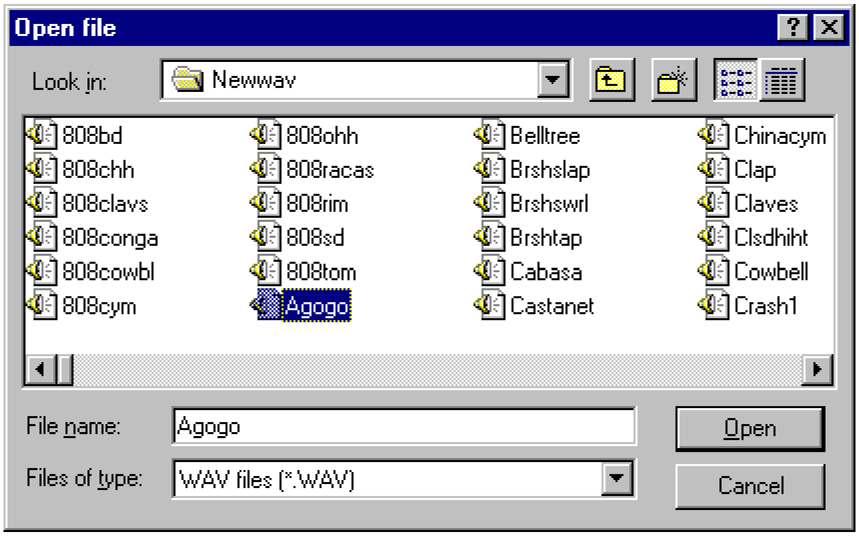

Select a file and click **Open** to load it.

Right-clicking the button brings up the Open settings dialog.

Acid WAV uses the Windows Audio Compression Manager (ACM) to decode encoded (usually compressed) WAV files. Any WAV format can be decoded as long as there is an ACM codec (coder / decoder) capable of handling that format installed on your system. Several codecs come with Windows. Others are installed by third party applications and add-ons such as Netscape Navigator and NetShow (e.g. Microsoft's MPEG Layer-3 codec).

You can get a complete list of all ACM codecs currently installed on your system by right-clicking the button in the toolbar. This brings up the Output format window (click its button to leave without actually changing the output format).

Saving

In order to save a file, click the button in the toolbar. This brings up a standard Windows Save dialog:

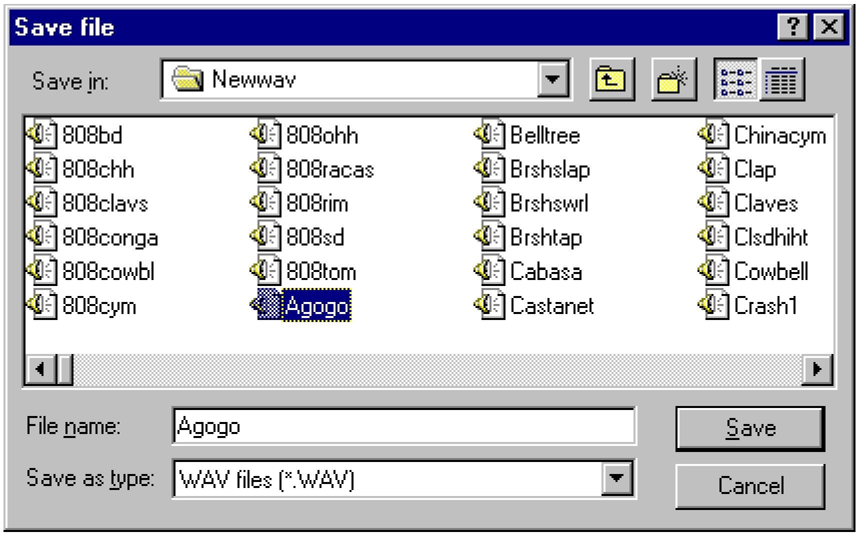

Enter a filename and click **Save** to write your sound file to disk.

Acid WAV remembers the format of loaded files and will use the same format when saving them (unless instructed to do otherwise by means of the Output format window). The default format (used e.g. when a file has been synthesized rather than loaded) is plain PCM, i.e. no encoding.

Both decoding and encoding are done transparently, i.e. they don't require any additional actions on your part.

Playing and recording

Playing

Playback is controlled from the toolbar. Click

to play the selected file section,

to pause and resume playback,

to stop playback.

If a start position has been marked, clicking will cause playback to start there. If an end position has also been marked, playback will eventually stop there. Once playback is underway, modifying the start and end positions has no effect on the file section being played unless the

button is clicked. Pausing and resuming playback with

causes the end position to be updated.

Recording

Recording is also controlled from the toolbar. Click

to start recording,

to stop recording.

If no file is loaded, clicking will cause the WAV properties window to be displayed. Set the desired file properties and click **OK** to start recording.

If a file is already loaded, recording works as follows:

- When neither start nor end positions are marked, the recorded sound is inserted ahead of the current data.

- When a start position but no end position is marked, the recorded sound is inserted at the start position.

- When both start and end positions are marked, the recorded sound overwrites the marked section.

Stereo recording can be limited to only one channel, leaving the other channel unaffected.

Right-click to bring up the Record settings window.

Analyzing sound data

Right-clicking the wave display when a file is loaded causes the analysis menu to pop up:

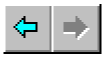

- Click **Statistics** for basic information about the selected file section.

- Click **Spectrum** for a frequency spectrum of the selected file section.

The last two menu entries are just like Spectrum, but only apply to one sound channel (Left or Right). They are only shown for stereo files.

Statistics

The **Statistics** window is invoked from the analysis pop-up menu. Its contents are largely self-explanatory.

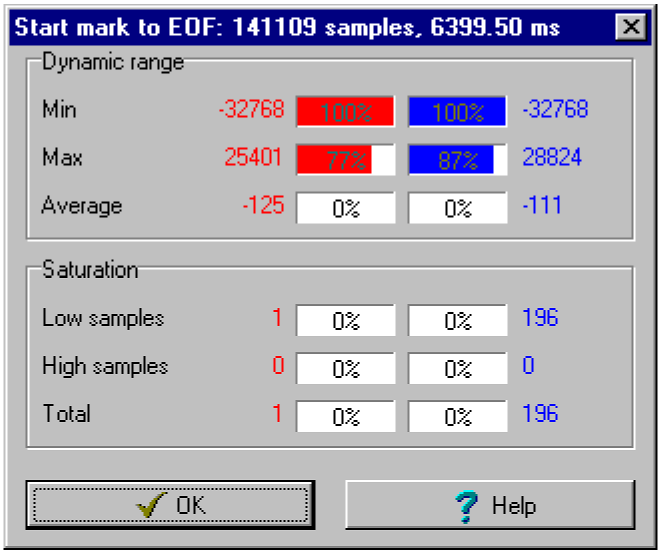

The title bar tells you which part of the file you are analyzing (whole file, start mark to End Of File or full selection - click here for details on selecting file sections) and its length.

The **Dynamic range** box summarizes the min, max and average sample values in each channel (left channel in red, right channel in blue). The available range for 16 bit files is -32768 to 32767, and the average value should be close to 0. For 8 bit files the available range is 0 to 255, and the average value should be close to 128.

An abnormal average value usually means that there is a DC component (i.e. a constant amplitude offset) in the sound. Such a DC component effectively reduces the available dynamic range and can cause various operations (e.g. mixing) to result in unnecessary distortion due to saturation. You can remove DC components using the Offset function in the Amplitude group.

The **Saturation** box tells you how many samples are saturated at the low and high ends of the available dynamic range (i.e. how many samples have the values -32768 and 32767 in 16 bit files, 0 and 255 in 8 bit files). A large number of saturated samples usually means audible distortion.

Spectrum

The **Spectrum** window is invoked from the analysis pop-up menu.

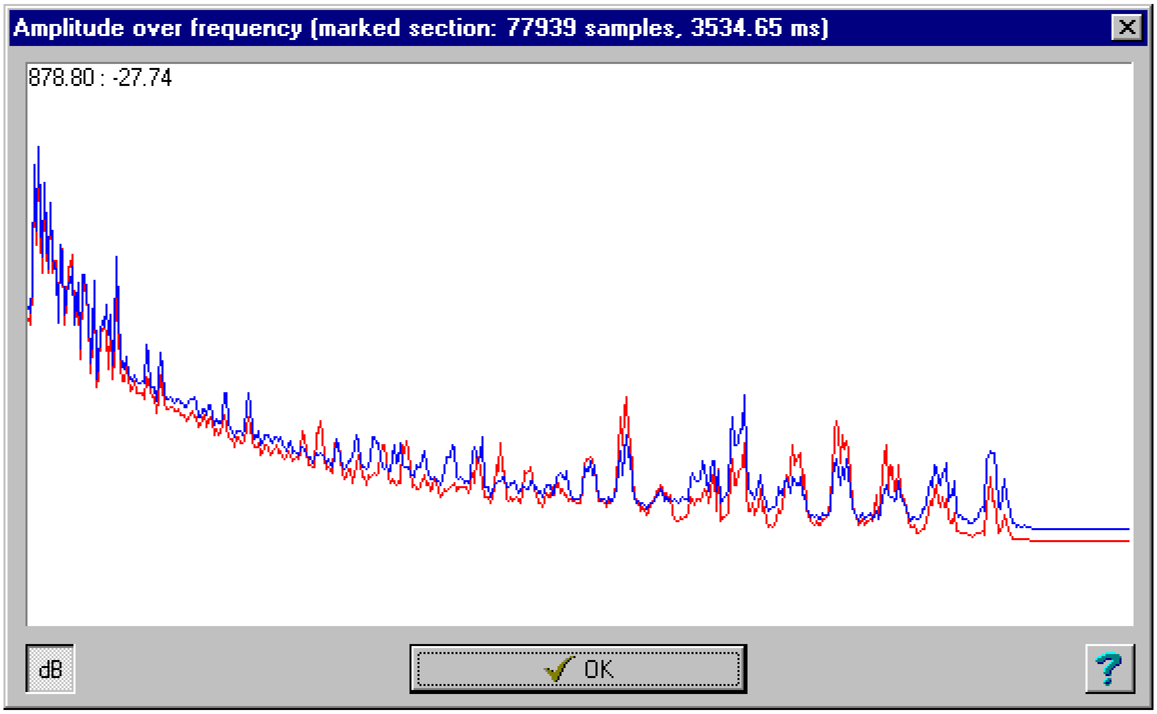

The title bar tells you which part of the file you are analyzing (whole file, start mark to End Of File or full selection - click here for details on selecting file sections) and its length.

The graph shows the spectrum, i.e. the amplitude as a function of frequency. As usual, the left channel is red and the right channel is blue. Frequency and amplitude values for the mouse cursor position are displayed in the upper left corner as "frequency (Hz) : amplitude".

If the **dB** button is depressed, the graph scale is logarithmic and amplitudes are expressed in decibel (with the full dynamic response supported by the sample size - 48 dB for 8 bits, 96 dB for 16 bits - serving as the reference level). Otherwise, the graph scale is linear and amplitudes are expressed as absolute sample values.

Basic editing

The **Basic** editing group contains the following functions:

- Trim
- Cut
- Copy
- Delete
- Paste
- Paste mix
- Insert file
- Mix files
- Insert silence
- Swap channels
- Duplicate channel
- Convert to mono / stereo
- Convert to 8 / 16 bits

See also: Changing file properties.

Trim

Trim is used to move alla data not in the selected section to the clipboard. In other words, it's the complement of Cut. It's typically used to isolate the selected section.

Right-click **Trim** to open its settings dialog:

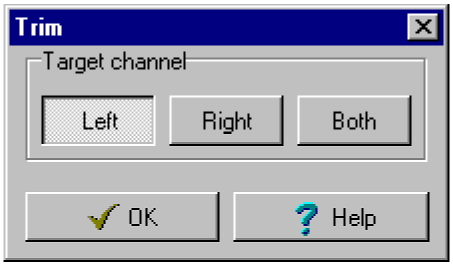

Clicking **Left** or **Right** causes data to be removed only from one channel. The resulting difference in duration is compensated by padding the selected channel with silence, leaving the overall file length unchanged. This setting does not affect what's written to the clipboard.

The **Target channel** setting is ignored when trimming monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Trim** will be executed.

Click 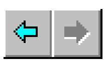 to abort.

Cut

Cut is used to move the selected section to the clipboard.

Right-click **Cut** to open its settings dialog:

Clicking **Left** or **Right** causes data to be removed only from one channel. The resulting difference in duration is compensated by padding the selected channel with silence, leaving the overall file length unchanged. This setting does not affect what's written to the clipboard.

The **Target channel** setting is ignored when cutting monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Cut** will be executed.

Click 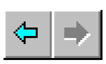 to abort.

Clicking the **Copy** button causes the selected file section to be copied to Acid WAV's clipboard.

Delete

Delete is used to remove the selected section. The difference between **Delete** and Cut is that no data is sent to the clipboard.

Right-click **Delete** to open its settings dialog:

Clicking **Left** or **Right** causes data to be removed only from one channel. The resulting difference in duration is compensated by padding the selected channel with silence, leaving the overall file length unchanged.

The **Target channel** setting is ignored when deleting data from monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Delete** will be executed.

Click 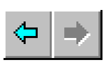 to abort.

Paste

Paste is used to copy data from the clipboard. The standard insertion rules apply.

Right-click **Paste** to open its settings dialog:

Clicking **Left** or **Right** causes data to be written only to one channel. The resulting difference in duration is compensated by padding the shorter channel with silence.

The **Target channel** setting is ignored when pasting to monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Paste** will be executed.

Click 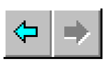 to abort.

Paste mix

Paste mix is used to mix the clipboard's contents into the file.

Right-click **Paste mix** to open its settings dialog:

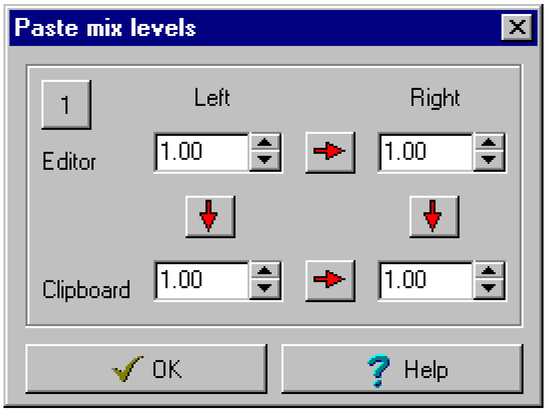

The edit boxes let you set amplification factors for the **Editor** and **Clipboard** data, individually for each channel. Negative values (causing phase inversion) are allowed.

Use the arrow buttons to copy amplification factors to neighbouring boxes.

The **Right** column is ignored when pasting data to monophonic files.

Clicking **1** causes all levels to be reset to 1.00, i.e. all channels are mixed down without any amplification.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Paste mix** will be executed.

Click 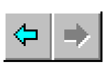 to abort and revert to the old settings.

Insert file

Insert file is used to copy data from a disk file. The standard insertion rules apply.

Right-click **Insert file** to open its settings dialog:

Clicking **Left** or **Right** causes data to be written to only one channel. The resulting difference in duration is compensated by padding the shorter channel with silence.

The **Target channel** setting is ignored when pasting to monophonic files.

Click 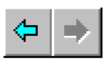 to abort.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. Either way, you will be presented with a standard Windows Open dialog:

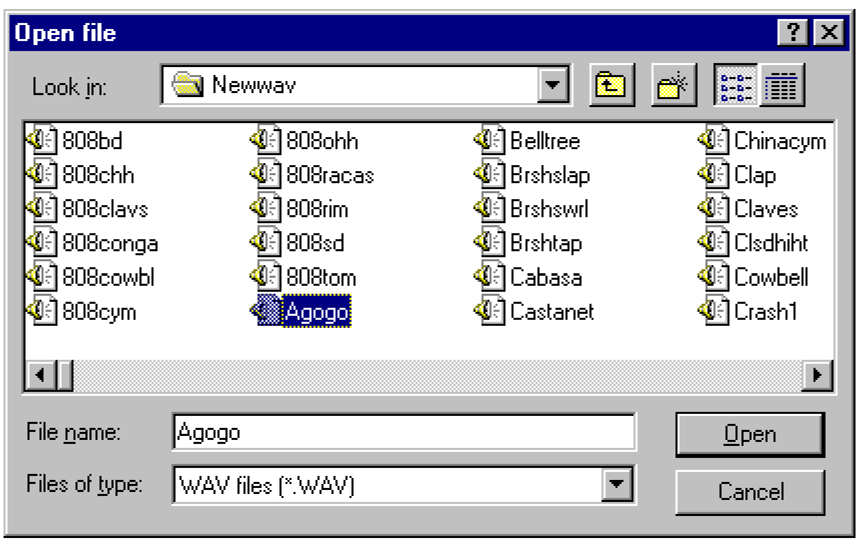

Select a file and click **Open** to insert it. Click **Cancel** or 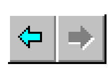 to abort.

Mix files

Mix files is used to mix one or more files into the editor.

Right-click **Mix files** to open its settings dialog:

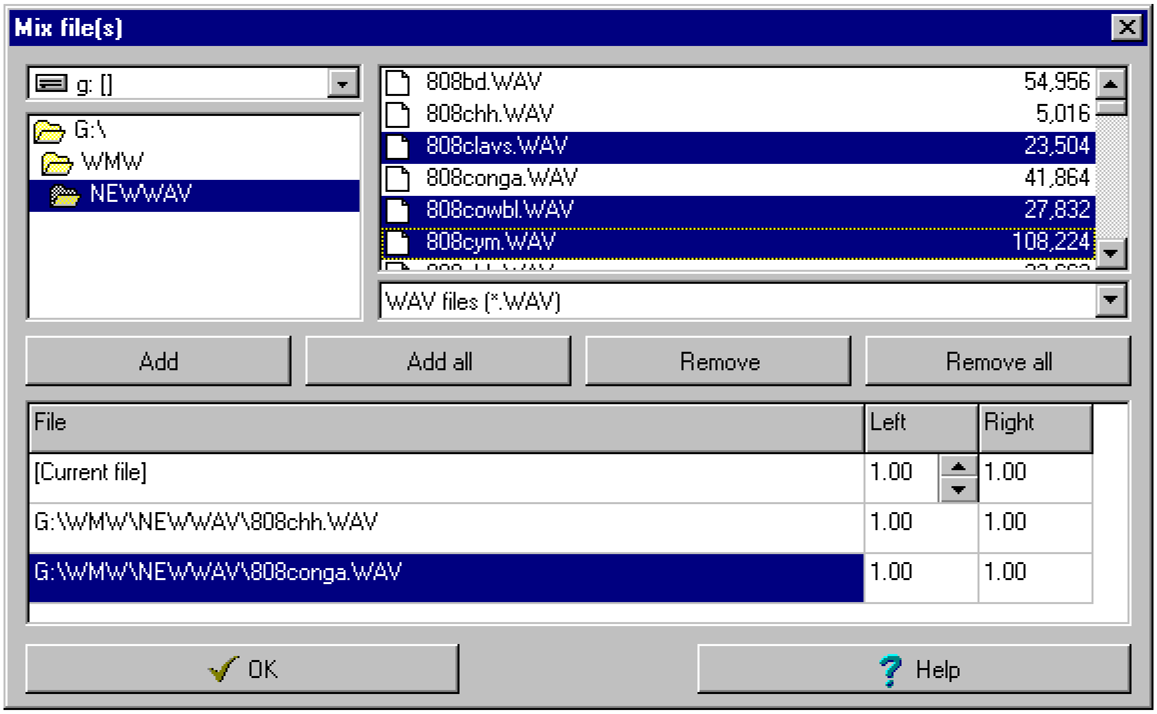

Select the files to mix using the controls in the upper part of the dialog window. Multiple file selections are supported (press the Shift key while clicking to select a range; press Ctrl while clicking to select additional individual files).

Use the buttons to **Add** the selected files or to **Add all** files in the current directory to the mix list in the lower part of the dialog window.

The mix list can be cleared using **Remove all**. Individual files can be selected and then **Remove**d.

The first entry in the list is the file currently being edited. It can not be removed.

Amplification factors are set in the **Left** and **Right** columns. Negative values (causing phase inversion) are allowed. All data is converted transparently to the target format (i.e. the format of the [Current file]) before being mixed down. If the target format is monophonic, the **Right** column is ignored.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Mix files** will be executed.

Click 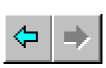 to abort and revert to the old settings.

Insert silence

Insert silence is used to add silent sections to files.

Right-click **Insert silence** to open its settings dialog:

If **Append to end** is depressed, the silent section is added after the last sample in the file, and the **Target channel** setting has no effect.

If **Append to end** is not depressed, the standard insertion rules apply.

When the dialog is called up from the main window, the **Duration** is expressed in the main window's time unit.

When the dialog is called up from the script editor, the **Duration** is expressed in milliseconds.

Click **OK** to accept the new settings If the dialog was called up from the main window, **Insert silence** will be executed.

Click 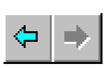 to abort and revert to the old settings.

Swap channels is used to make stereo channels change place in the selected section. It can't be applied to monophonic files.

Duplicate channel

Duplicate channel is used to copy one stereo channel in the selected section to the other channel. It can't be applied to monophonic files.

Right-click **Duplicate channel** to open its settings dialog:

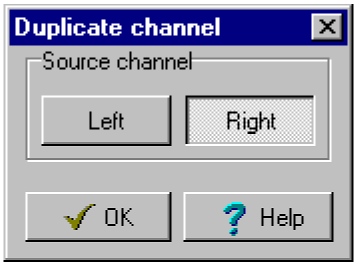

Click **OK** to confirm the current setting or either button in the **Target channel** box to update it. If the dialog was called up from the main window, **Duplicate channel** will be executed.

Click 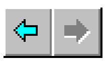 to abort.

When a monophonic file is being edited, the next to last button in the **Basic** group is labeled **Convert to stereo**. Clicking it causes all sound data to be duplicated into a second sound channel, creating a stereophonic file.

When a stereophonic file is being edited, the button is labeled **Convert to mono**. Clicking it causes all sound data in the two channels to be averaged into a single channel, creating a monophonic file.

Any file section selection is ignored (but preserved).

When a 16 bit file is being edited, the last button in the **Basic** group is labeled **Convert to 8 bits**. Clicking it causes all samples to be reduced to 8 bit resolution, preserving their amplitude relative to the available dynamic range. Note that this operation causes information to be lost, introducing noise.

When an 8 bit file is being edited, the button is labeled **Convert to 16 bits**. Clicking it causes all samples to be scaled up to 16 bit resolution, preserving their amplitude relative to the available dynamic range. No information is lost (and no noise added) in this process.

Any file section selection is ignored (but preserved).

Changing file properties

The fastest way to change the word size (8 or 16 bits) or the number of channels (1 or 2) of a file is to use the dedicated **Basic** editing functions.

An alternative method allowing you to change sampling rate, word size and number of channels in one shot is to use the Copy and Paste mix functions:

- Click **Copy** to send the selected file section to the clipboard.

- Click 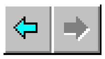 in the toolbar to clear the editor. Alternatively, click 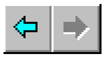 and switch to the new Acid WAV window.

- Click **Insert silence**. Since no file is loaded, this brings up the WAV properties window. Set the desired properties and click **OK**. A new, silent file will be created.

  (You may want to right-click **Insert silence** instead of left-clicking it. This allows you to change the duration of the new file to something shorter than the clipboard contents, saving you any trimming of the final result later on.)

- Click **Paste mix** (right-click if you are not sure that the amplification levels are set appropriately). Done!

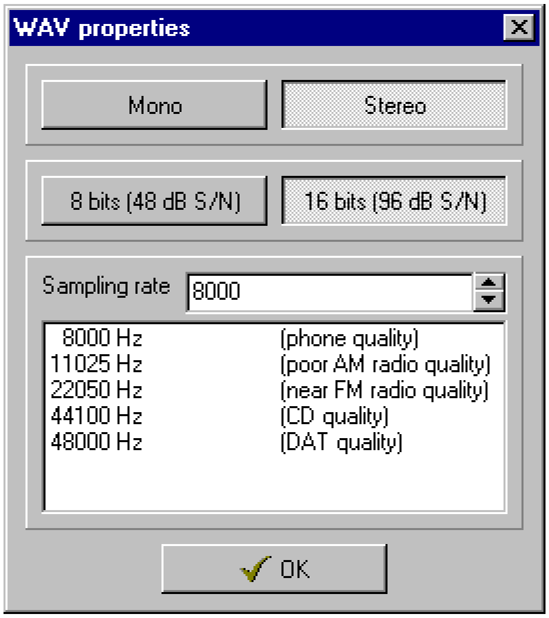

The **WAV properties** window is displayed by Acid WAV when it needs to create a new file from scratch, e.g. for synthesis or recording. After selecting the desired properties, click **OK** to continue or 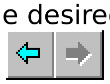 to discard all changes and cancel the operation.

Acid WAV's clipboard is a file (not the Windows clipboard) used for temporary storage of sound data. You can send data to the clipboard with Trim, Cut or Copy. All Acid WAV windows share the same clipboard (unless you have several copies of Acid WAV installed in different directories, in which case each copy will maintain its own clipboard).

Writing new data to the clipboard causes any data previously sent there to be lost. The clipboard is also cleared when the last Acid WAV window is closed.

The clipboard's contents can be pasted into the editor using Paste and Paste mix. In addition, Ring modulate can use the clipboard as its modulator source, and Reduce noise requires the clipboard to contain samples of the noise to be removed. Sound data which was written to the clipboard in a format different than that of the destination (e.g. 16 bit data being pasted into an 8 bit file) is reformatted transparently by Acid WAV.

See also: Selecting file sections.

No selection

When neither start nor end positions are marked in the main window's wave display, new data is inserted ahead of the first sample in the file.

No end position

When only a start position is marked in the mani window's wave display, new data is inserted at the start position.

Full selection

When both start and end positions are marked in the main window's wave display, new data replaces the highlighted section.

Stereo files

Differences in duration between stereo channels caused by inserting (or overwriting) data into only one channel are compensated by padding the shorter channel with silence.

Advanced editing

Advanced editing is a collective term for anything done with the following function groups:

- Time functions
- Amplitude functions
- Frequency functions

Time functions

The **Time** editing group contains the following functions:

- Echo
- Hall reverb
- Chapel reverb
- Chamber reverb
- Custom reverb
- Virtual room
- Reverse
- Stretch
- Tremolo
- Vibrato
- Chorus
- Flange
- Ring modulate

Echo

Echo is used to create a basic echo effect throughout the selected section. Use Virtual room for more realistic room simulations.

Right-click **Echo** to open its settings dialog:

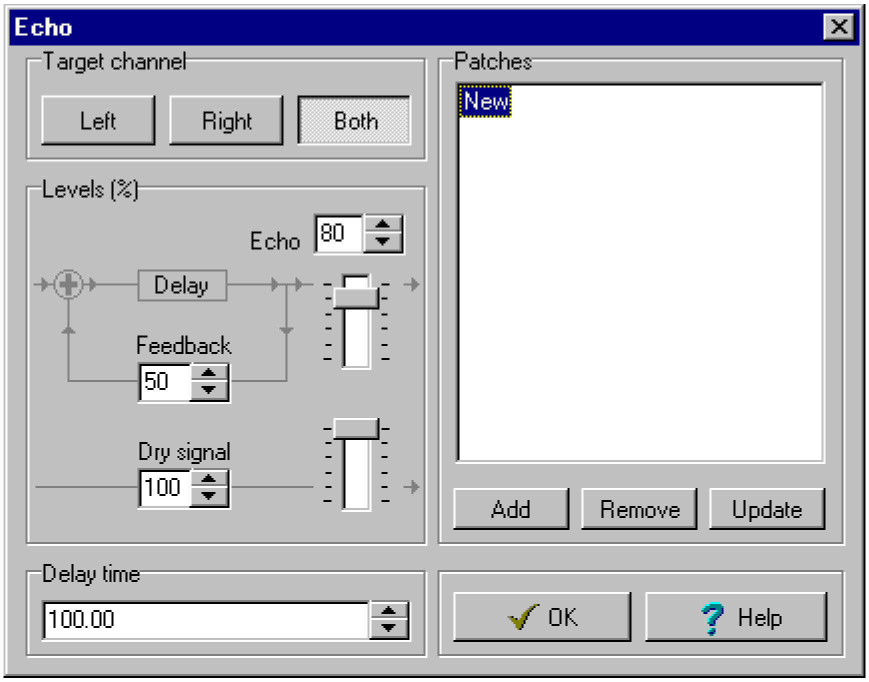

You have the following parameters at your disposal:

- **Echo** is the effect mix level. At 100%, the first echo will have the same amplitude as the original sound.
- **Feedback** is the amplitude ratio between successive echoes. At 50%, the amplitude will fall by half for each new echo.
- **Dry signal** is the mix level of the original sound.
- **Delay time** is the separation between the original sound and the first echo, as well as between successive echoes.

When the dialog is called up from the main window, the delay time is expressed in the main window's time unit.

When the dialog is called up from the script editor, the delay time is expressed in milliseconds.

Selecting **Left** or **Right** causes only one channel to be echoed. When **Both** is selected, the left and right channel are treated separately.

Creating echoes in real rooms necessarily implies mixing stereo channels. You can therefore improve the realism of the effect by partially mixing channels or (even better) by reducing the original file to mono, echoing it and mixing the result into the echoed stereo file. This is a good candidate for a simple script.

You can save all settings to (and restore them from) your own echo Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Echo** will be executed.

Click 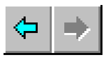 to abort and revert to the old settings.

Hall reverb

Hall reverb is used to apply a large-room reverb algorithm to the selected section.

See Chapel reverb and Chamber reverb for other special reverb algorithms, Custom reverb for a generic algorithm and Virtual room for a realistic room simulator.

Right-click **Hall reverb** to open its settings dialog:

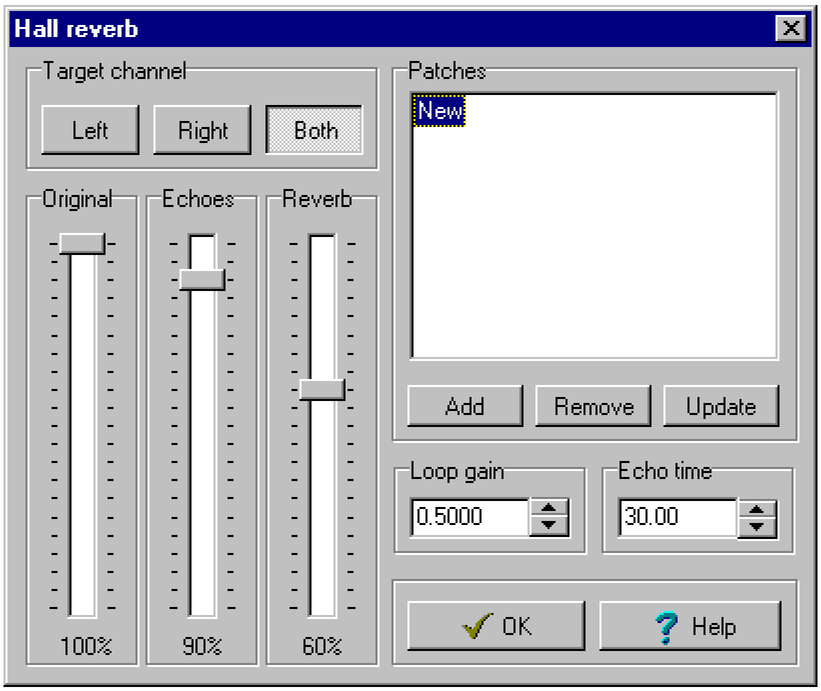

Like most synthetic reverbs, **Hall reverb** creates its output by mixing three signals:

- the **Original** sound,
- early **Echoes** from a multitap delay line (used to simulate early, distinct reflections from nearby walls) and
- the diffuse **Reverb** proper, created by looping the sound through a network of delay lines (used to simulate mixed, multiple reflections from all walls).

Apart from the mix levels for these sources, you can set the

- **Echo time** for the early reflection delay line (use longer echo times to create the impression of a larger room) and the
- **Loop gain** for the diffuse reverb section. This determines how fast the reverb will decay. Use larger gains to make the reverberations last longer, creating the impression of a more complicated room geometry (more walls).

When the dialog is called up from the main window, the echo time is expressed in the main window's time unit.

When the dialog is called up from the script editor, the echo time is expressed in milliseconds.

Selecting **Left** or **Right** causes only one channel to be reverberated. When **Both** is selected, the left and right channel are treated separately.

In real rooms, reverberation necessarily implies mixing between stereo channels. You can therefore improve the realism of the effect by partially mixing channels or (even better) by reducing the original file to mono, reverberating it and mixing the result into the reverberated stereo file. This is a good candidate for a simple script.

The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own hall reverb Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Hall reverb** will be executed.

Click 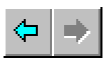 to abort and revert to the old settings.

Chapel reverb

Chapel reverb is used to apply a medium-sized-room reverb algorithm to the selected section.

See Hall reverb and Chamber reverb for other special reverb algorithms, Custom reverb for a generic algorithm and Virtual room for a realistic room simulator.

Right-click **Chapel reverb** to open its settings dialog:

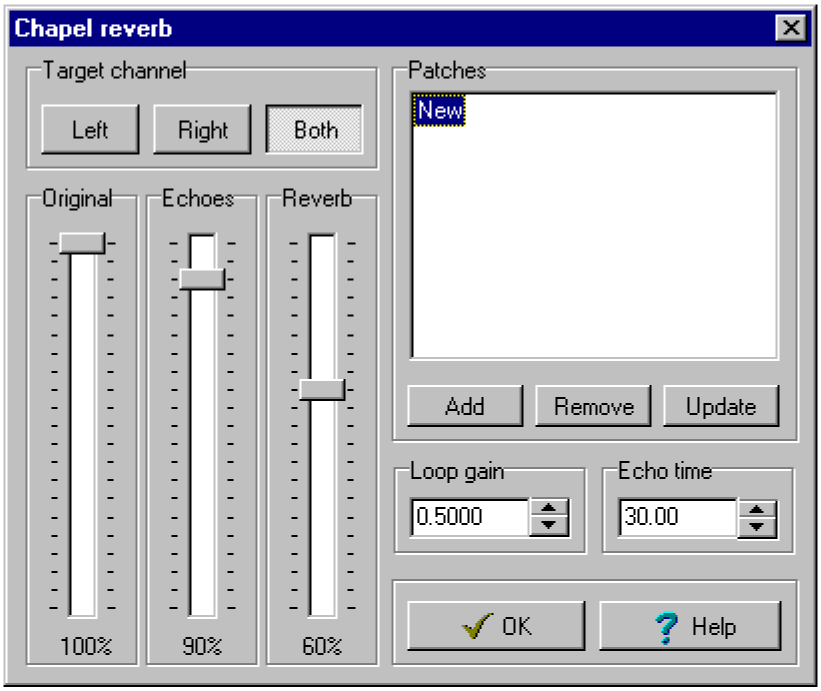

Like most synthetic reverbs, **Chapel reverb** creates its output by mixing three signals:

- the **Original** sound,
- early **Echoes** from a multitap delay line (used to simulate early, distinct reflections from nearby walls) and
- the diffuse **Reverb** proper, created by looping the sound through a network of delay lines (used to simulate mixed, multiple reflections from all walls).

Apart from the mix levels for these sources, you can set the

- **Echo time** for the early reflection delay line (use longer echo times to create the impression of a larger room) and the
- **Loop gain** for the diffuse reverb section. This determines how fast the reverb will decay. Use larger gains to make the reverberations last longer, creating the impression of a more complicated room geometry (more walls).

When the dialog is called up from the main window, the echo time is expressed in the main window's time unit.

When the dialog is called up from the script editor, the echo time is expressed in milliseconds.

Selecting **Left** or **Right** causes only one channel to be reverberated. When **Both** is selected, the left and right channel are treated separately.

In real rooms, reverberation necessarily implies mixing between stereo channels. You can therefore improve the realism of the effect by partially mixing channels or (even better) by reducing the original file to mono, reverberating it and mixing the result into the reverberated stereo file. This is a good candidate for a simple script.

The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own chapel reverb Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Chapel reverb** will be executed.

Click 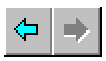 to abort and revert to the old settings.

Chamber reverb

Chamber reverb is used to apply a small-room reverb algorithm to the selected section.

See Hall reverb and Chapel reverb for other special reverb algorithms, Custom reverb for a generic algorithm and Virtual room for a realistic room simulator.

Right-click **Chamber reverb** to open its settings dialog:

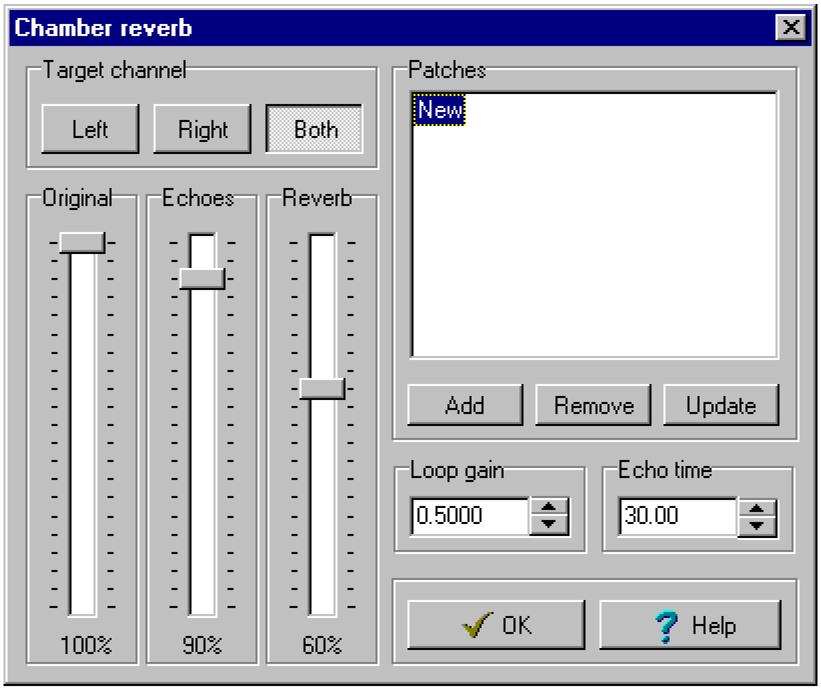

Like most synthetic reverbs, **Chamber reverb** creates its output by mixing three signals:

- the **Original** sound,
- early **Echoes** from a multitap delay line (used to simulate early, distinct reflections from nearby walls) and
- the diffuse **Reverb** proper, created by looping the sound through a network of delay lines (used to simulate mixed, multiple reflections from all walls).

Apart from the mix levels for these sources, you can set the

- **Echo time** for the early reflection delay line (use longer echo times to create the impression of a larger room) and the
- **Loop gain** for the diffuse reverb section. This determines how fast the reverb will decay. Use larger gains to make the reverberations last longer, creating the impression of a more complicated room geometry (more walls).

When the dialog is called up from the main window, the echo time is expressed in the main window's time unit.

When the dialog is called up from the script editor, the echo time is expressed in milliseconds.

Selecting **Left** or **Right** causes only one channel to be reverberated. When **Both** is selected, the left and right channel are treated separately.

In real rooms, reverberation necessarily implies mixing between stereo channels. You can therefore improve the realism of the effect by partially mixing channels or (even better) by reducing the original file to mono, reverberating it and mixing the result into the reverberated stereo file. This is a good candidate for a simple script.

The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own chamber reverb Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Chamber reverb** will be executed.

Click 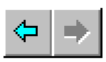 to abort and revert to the old settings.

Custom reverb

Custom reverb is used to apply a generic reverb algorithm to the selected section.

See Hall reverb, Chapel reverb and Chamber reverb for three special reverb algorithms and Virtual room for a realistic room simulator.

Right-click **Custom reverb** to open its settings dialog:

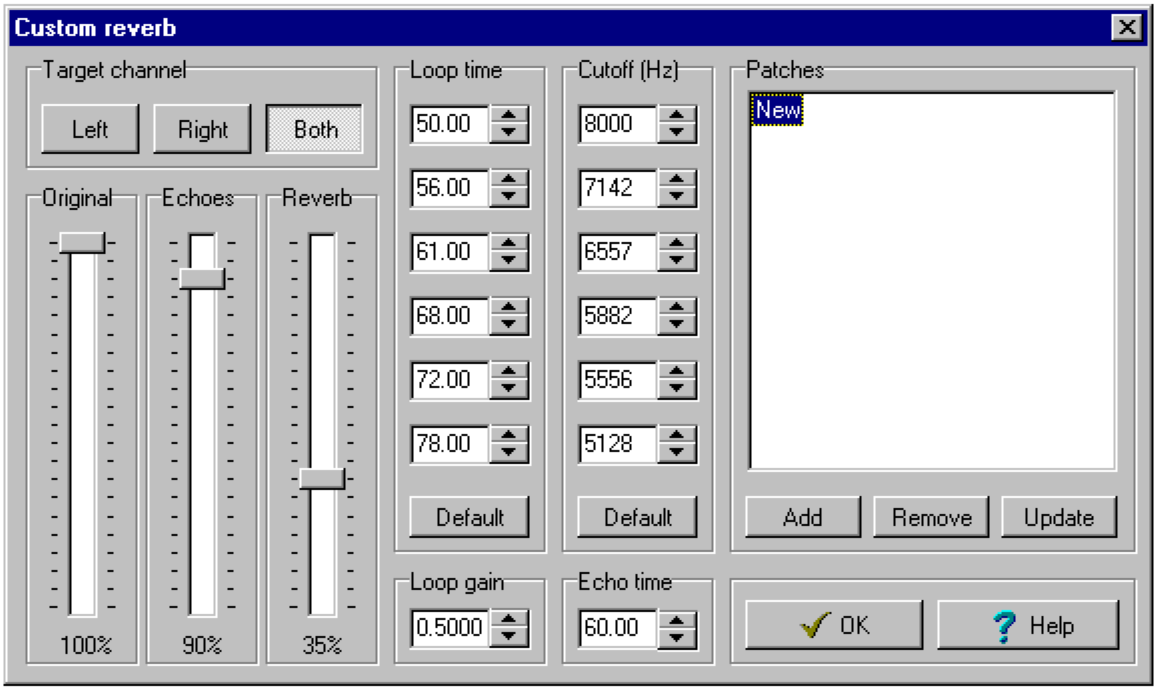

Like most synthetic reverbs, **Custom reverb** creates its output by mixing three signals:

- the **Original** sound,
- early **Echoes** from a multitap delay line (used to simulate early, distinct reflections from nearby walls) and
- the diffuse **Reverb** proper, created by looping the sound through a network of delay lines (used to simulate mixed, multiple reflections from all walls).

As with the special reverb algorithms, you can set the mix levels for these sources as well as the

- **Echo time** for the early reflection delay line (use longer echo times to create the impression of a larger room) and the
- overall **Loop gain** for the diffuse reverb section. This determines (in part - see below) how fast the reverb will decay. Use larger gains to make the reverberations last longer, creating the impression of a more complicated room geometry (more walls).

You also have a handle on the inner workings of the reverb section through the **Loop time** and **Cutoff** controls. This section is made up of six smaller reverberator units working in parallel. Each one has its own delay loop containing an amplifier and a simple low pass filter.

The amplifier simulates overall attenuation as the sound travels through air and bounces off walls; the low pass filter makes the simulation more realistic by attenuating high frequencies faster than low ones.

Making a **Loop time** longer is like moving two walls further apart; lowering the **Cutoff** frequency is like making the walls "softer" and the air between them more humid, resulting in increased absorption of high frequencies over low ones.

Loop gain and **Loop time** are of primary importance in determining the overall decay time. A shorter **Loop time** will counteract a larger **Loop gain**, but by packing more "bounces" per time unit it will also contribute to a smoother reverb. The effect of the **Cutoff** settings becomes successively more noticeable over longer decay times, and can become a dominating feature of the sound picture after a few seconds.

When the dialog is called up from the main window, times are expressed in the main window's time unit.

When the dialog is called up from the script editor, times are expressed in milliseconds.

Selecting **Left** or **Right** causes only one channel to be reverberated. When **Both** is selected, the left and right channel are treated separately.

In real rooms, reverberation necessarily implies mixing between stereo channels. You can therefore improve the realism of the effect by partially mixing channels or (even better) by reducing the original file to mono, reverberating it and mixing the result into the reverberated stereo file. This is a good candidate for a simple script.

The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own custom reverb Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Custom reverb** will be executed.

Click 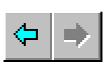 to abort and revert to the old settings.

Virtual room

Virtual room is used to run the selected section through a realistic room acoustics simulator.

See Echo, Hall reverb, Chapel reverb, Chamber reverb and Custom reverb for synthetic echo and reverb effects.

Right-click **Virtual room** to open its settings dialog:

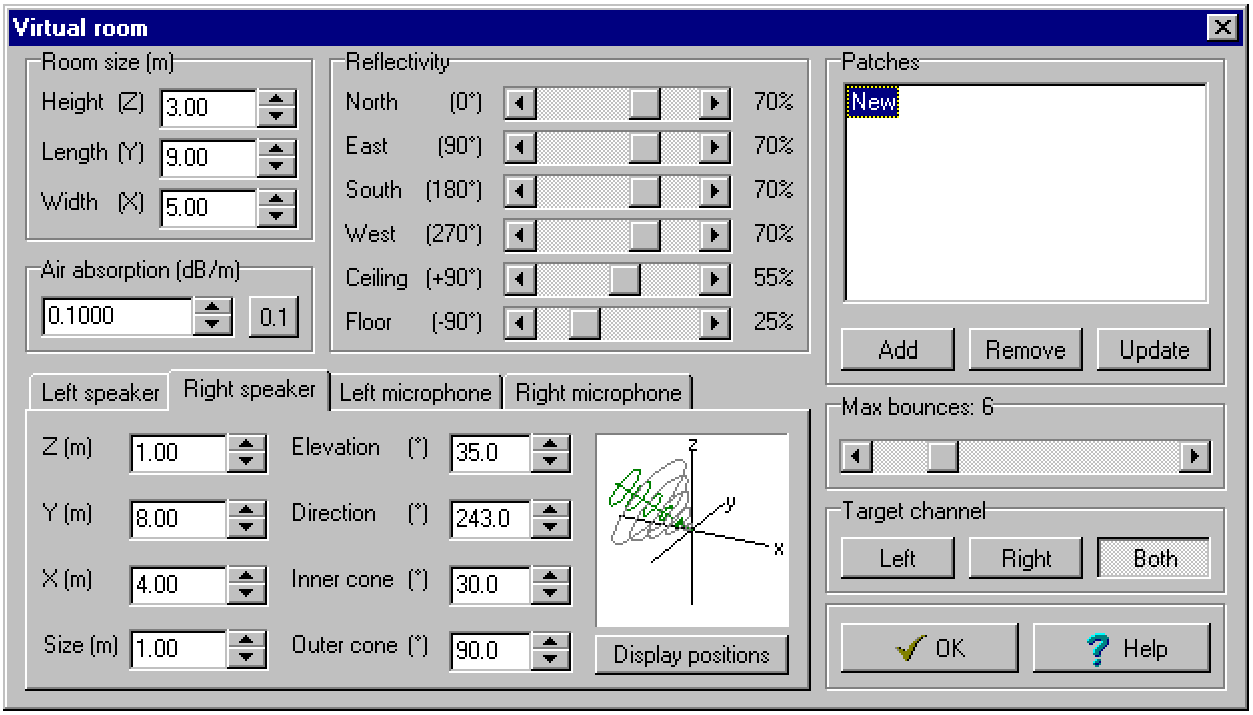

Setting up a room acoustics simulation involves three steps (treated in detail below):

- Specifying room properties (geometry and absorption coefficients).
- Specifying speaker and microphone properties (positions and directional characteristics).
- Specifying the degree of approximation to use in computations.

You can save all settings to (and restore them from) your own virtual room Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Virtual room** will be executed.

Click 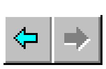 to abort and revert to the old settings.

Room properties

Room properties are entered in the upper portion of the **Virtual room** dialog. Apart from the **Room size**, you can vary the **Air absorption** and the **Reflectivity** of walls, floor and ceiling.

Air absorption may require some explaining. Sound coming from a distant source is weaker than sound coming from an identical but closer source because of two effects: the dissipation of sound energy into heat (caused by air being a less than ideal medium) and the geometrical "dilution" of sound energy over a growing wavefront area. The **Air absorption** parameter is used to control dissipation. Setting it to zero gives you a room filled with an ideal, non-dissipative medium, but it won't change the fact that wavefront areas grow with distance.

By the way, in case you're still using body parts to measure distances, 1 meter = 3.3 feet.

Reflectivity is the fraction of the sound's amplitude that survives a bounce from a surface. 100% means that the surface is a perfect reflector, 0% that it's a perfect absorber generating no echoes (this is what's usually strived for in "silent rooms", recording studios and the like).

In real life, neither air absorption nor surface reflectivity are independent of the sound's frequency. If you don't mind some heavy duty computations, you can simulate the effects of frequency-dependent absorption and reflecitivity by splitting the source signal into adjacent frequency bands (e.g. with Band pass or Equalize), running each band through a virtual room with its own air absorption and reflection coefficients, and finally mixing the results back together. This is a good candidate for a script.

Speakers and microphones

Each simulated speaker has a position (**X**, **Y** and **Z** coordinates) a **Size**, an orientation (**Elevation** and **Direction**) and directionality characteristics (**Inner cone** and **Outer cone**).

- **Size** is the radius of a sphere from which the source signal expands into the virtual room. This parameter sets the scale for the amplitude attenutation over distance caused by the "dilution" of sound energy over a growing wavefront area. A larger **Size** (initial wavefront) results in slower attenuation.

- **Elevation** and **Direction** are angles best understood by observing how they affect the directional display:

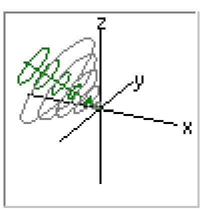

- By definition, a speaker radiating sound with the same intensity in all directions has no orientation. Elevation and direction are made meaningful by the **Inner cone** and **Outer cone** parameters:

  The **Inner cone** (outlined by a series of green circles in the directional display) is the region throughout which the speaker emits sound with full intensity.

  The **Outer cone** (outlined by a series of gray circles in the directional display) is a region throughout which the emitted intensity falls off from full (at the surface of the inner cone) to zero (at the surface of the outer cone). The transition is linear in amplitude, quadratic in intensity.

The speaker does not radiate outside of the outer cone.

Microphones are similar to speakers. The main difference is that they are pointlike (their **Size** is always 0). Their **Inner cone** and **Outer cone** parameters control sensitivity: zero outside the outer cone, 100% inside the inner cone, linear in amplitude between the two.

Microphones and speakers are also affected by the **Target channel** setting. Selecting **Left** or **Right** causes only one channel to be routed through the virtual room. When **Both** is selected, the left and right channel are treated jointly, i.e. sound emitted by the left speaker is also picked up by the right microphone and vice versa, like in a real room.

The **Target channel** setting is ignored for monophonic files, which are always routed through the **Left** microphone and speaker.

Visualizing positions and orientations from numbers alone can be difficult. Click Display positions to bring up a wireframe display of the room:

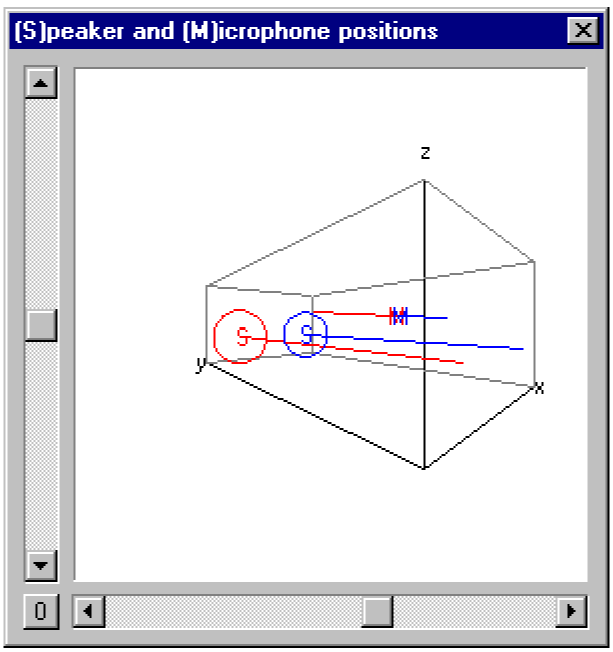

Speakers are marked with **S**s, microphones with **M**s. Orientations are shown as lines pointing away from microphones and speakers. The usual color codes apply: left is red and right is blue. Use the scroll bars to rotate the display. Click **0** to bring it back to the default position.

Note that the room display is modeless, so you can keep it open while you edit room properties in the **Virtual room** dialog.

Degree of approximation

Acid WAV simulates room acoustics by adding up the effects of successive echoes: the original sound hits the walls (and the floor, and the ceiling) and bounces off, generating new wavefronts. Each new wavefront in turn goes on to bounce off the other walls (and the floor, and the ceiling) generating even more, "second generation" bounced wavefronts. This can in principle go on forever, with the number of wavefronts growing exponentially.

Two things save the day:

- The attenuation as sound travels through air and bounces off walls. As energy is dissipated, the contribution of new wavefronts to the sound drops off, and we can get away with neglecting bounces beyond a certain "generation".

- Each bounce takes time to happen (sound travels through air at roughly 340 m/s), and we are only interested in a limited time frame.

The number of bounces to keep track of is set in the **Max bounces** box. 0 bounces means "simulate no echoes, just tell me what the direct signal sounds like after it's traveled from the speakers to the microphones". 1 bounce means "simulate first generation bounced wavefronts" (six of them, plus the original signal) and so on.

Reverse

Reverse is used to time-invert the selected section.

Right-click **Reverse** to open its settings dialog:

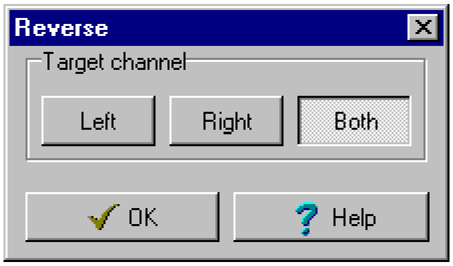

Clicking **Left** or **Right** causes only one channel to be reversed.

The **Target channel** setting is ignored when reversing monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Reverse** will be executed.

Click 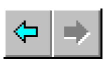 to abort.

Stretch

Stretch is used to change the duration of the selected section without changing its pitch.

Acid WAV's stretcher uses an advanced, Fourier-based algorithm primarily aimed at samples of music instruments containing a dominating frequency (or set of harmonic frequencies).

Right-click **Stretch** to open its settings dialog:

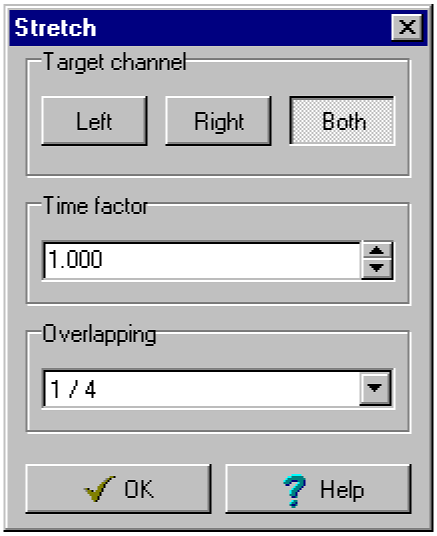

Use the **Time factor** box to set the amount of stretching. Values between 0.5 (causing the duration to be reduced by half) and 2.0 (causing the duration to be doubled) are allowed.

The **Overlapping** parameter determines the amount of cross-fading between successive sound chunks (like all functions, the stretcher reads and writes sound data in chunks a few kB in length). A setting of 1 / 1 means that the first half of each chunk is cross-faded with the second half of the previous chunk, making all chunks (except for the first and the last one) completely cross-faded with other chunks. A setting of 1 / 2 means that the first and the last first quarter of each chunk is cross-faded with other chunks, leaving the central half untouched, and so on. More overlapping tends to improve the result by smoothing out the transition between chunks. It also increases the computation time.

Selecting **Left** or **Right** causes only one channel to be stretched. The resulting difference in duration is compensated by padding the shorter channel with silence, leaving the overall file length unchanged.

The **Target channel** setting is ignored for monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Stretch** will be executed.

Click 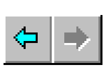 to abort and revert to the old settings.

Tremolo

Tremolo is used to impose a periodic amplitude envelope on the selected section. Use Envelop for more advanced amplitude shaping.

See also Vibrato.

Right-click **Tremolo** to open its settings dialog:

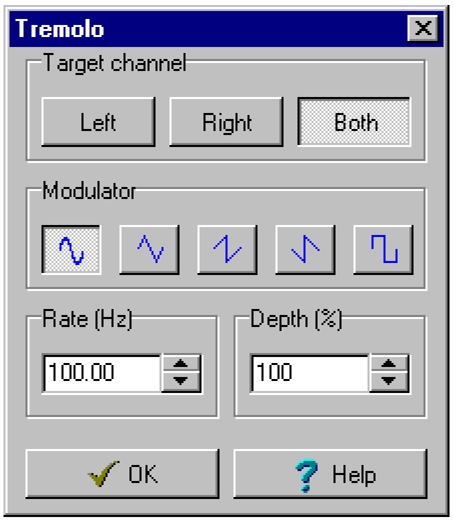

The **Modulator** box lets you choose between the following waveshapes:

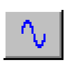 Sine wave.

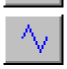 Triangle wave.

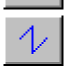 Up ramp.

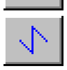 Down ramp.

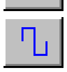 Square wave.

Set the number of cycles per second in the **Rate** box and the amount of modulation in the **Depth** box. A modulation depth of 100% causes the amplitude factor to vary between 1 and 0; a depth of 50% causes the amplitude factor to vary between 1 and 0.5.

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Tremolo** will be executed.

Click 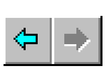 to abort and revert to the old settings.

Vibrato

Vibrato is used to impose a periodic frequency envelope on the selected section. Use Modulate for more advanced frequency modulation.

See also Tremolo.

Right-click **Vibrato** to open its settings dialog:

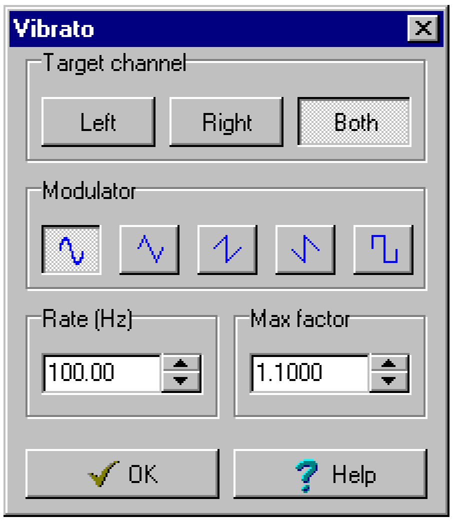

The **Modulator** box lets you choose between the following waveshapes:

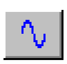 Sine wave.

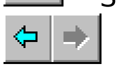 Triangle wave.

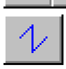 Up ramp.

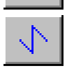 Down ramp.

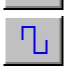 Square wave.

Set the number of cycles per second in the **Rate** box and the amount of modulation in the **Max factor** box. A max factor of 1.1 causes the frequency factor to vary between 1.1 and 1 / 1.1; a max factor of 2.0 causes the frequency factor to vary between 2.0 and 1 / 2.0.

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Vibrato** will be executed.

Click 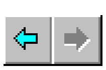 to abort and revert to the old settings.

Chorus

Chorus is used to mix the selected section with a delayed version of itself, with the delay undergoing a periodic variation to create the interference effects typical of choirs.

Right-click **Chorus** to open its settings dialog:

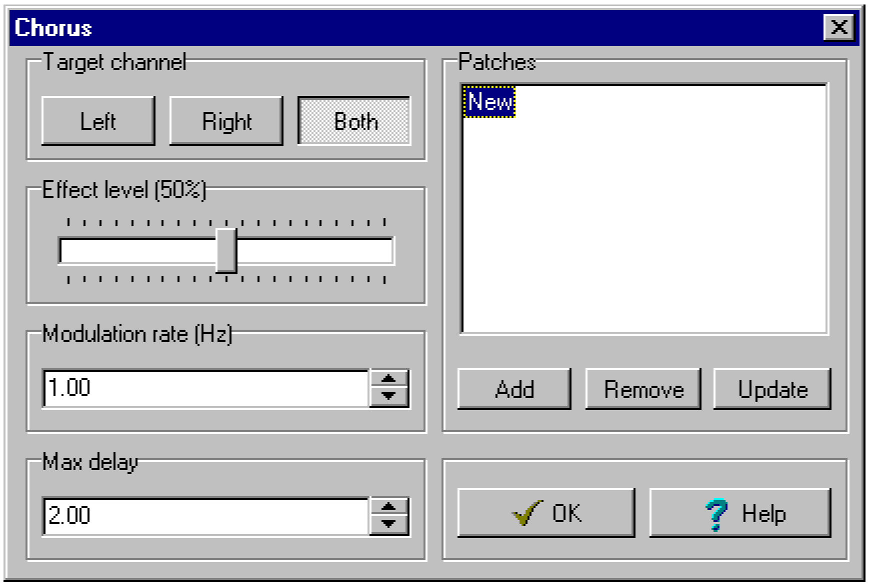

The **Effect level** control is used to set the amplitude of the delayed signal relative to the original signal.

The **Modulation rate** is the number of times that the delay is swept from 0 to **Max delay** and back in one second.

When the dialog is called up from the main window, the **Max delay** time is expressed in the main window's time unit.

When the dialog is called up from the script editor, the **Max delay** time is expressed in milliseconds.

For realistic chorus effects, the **Max delay** and the **Modulation rate** should be kept fairly small (a few milliseconds and Hertz, respectively).

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own chorus Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Chorus** will be executed.

Click 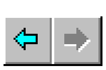 to abort and revert to the old settings.

Flange

Flange is a cross between Echo and Chorus on steroids. It's used to mix the selected section with an echoed version of itself, with the echo's delay undergoing a periodic variation to create time-varying interference effects.

Right-click **Flange** to open its settings dialog:

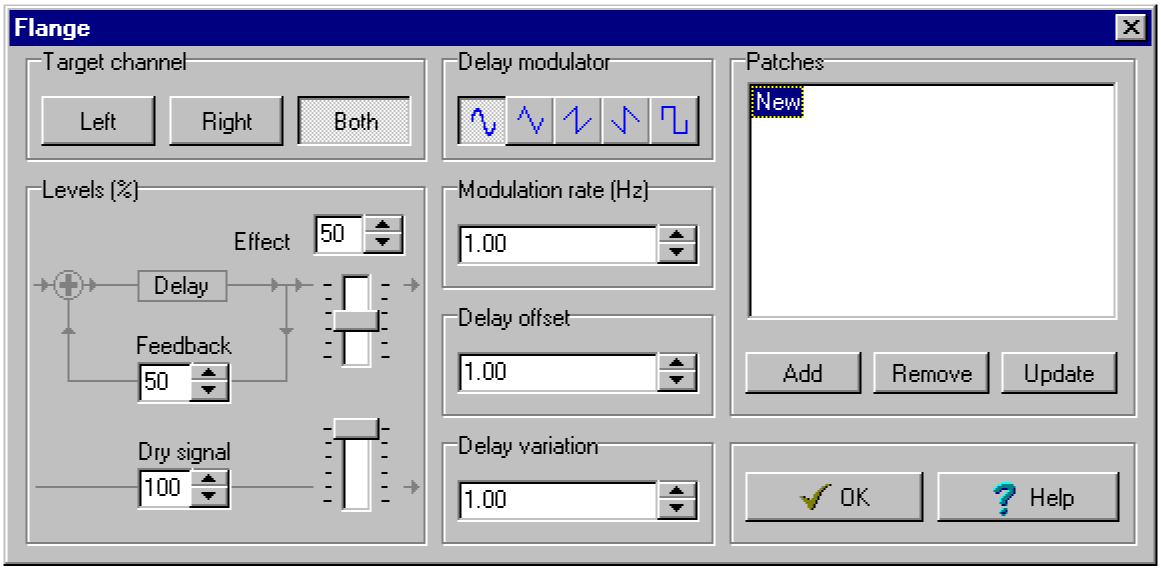

You have the following parameters at your disposal:

- **Effect** is the mix level of the echoed sound.
- **Feedback** is the amplitude ratio between successive echoes. At 50%, the amplitude will fall by half for each new echo.

  The non-linearity introduced by the feedback loop in the echo section is the most important difference between **Chorus** and **Flange**. With large feedback values, **Flange** can create truly warped sounds. With more moderate settings, it's one of the most popular effects among electric guitarists and the core of "grunge" sound.
- **Dry signal** is the mix level of the original sound.
- **Delay modulator** waveshapes:

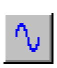 Sine wave (choose this modulator and set **Feedback** = 0 to reduce **Flange** to **Chorus**).

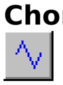 Triangle wave.

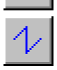 Up ramp.

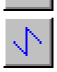 Down ramp.

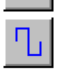 Square wave.

- **Modulation rate** is the number of times that the delay in the echo section is swept from **Delay offset** - **Delay variation** to **Delay offset** + **Delay variation** and back in one second.

The most interesting flange effects are created with fairly small modulation rates and delay times (a few Hertz and milliseconds, respectively).

When the dialog is called up from the main window, **Delay offset** and **Delay variation** times are expressed in the main window's time unit.

When the dialog is called up from the script editor, **Delay offset** and **Delay variation** times are expressed in milliseconds.

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own flanger Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Flange** will be executed.

Click 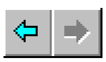 to abort and revert to the old settings.

Ring modulate

Ring modulate is used to multiply the selected section with a periodic signal.

A common use is the creation of "robot" voices and other synthetic-sounding effects.

Right-click **Ring modulate** to open its settings dialog:

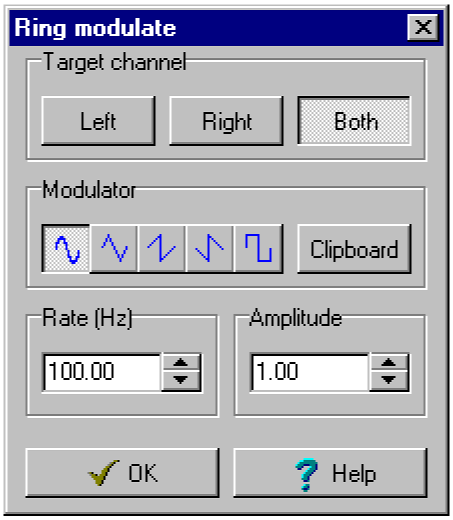

The **Modulator** box lets you choose between the following signal sources:

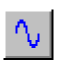 Sine wave.

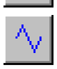 Triangle wave.

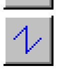 Up ramp.

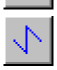 Down ramp.

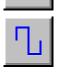 Square wave.

The clipboard. This option allows you to use any synthetic or sampled signal as a modulator.

Set the number of cycles per second in the **Rate** box. This setting is ignored when the clipboard is the modulator.

Use the **Amplitude** box to set the amplification factor to be applied to the modulator.

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Ring modulate** will be executed.

Click 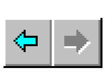 to abort and revert to the old settings.

Amplitude functions

The **Amplitude** editing group contains the following functions:

- Amplify
- Envelop
- Maximize
- Offset
- Invert
- Gate
- Clip
- Distort
- Rectify
- Quantize
- Interpolate
- Gated delete
- Compress / Expand

Amplify

Amplify is used to modify sound levels in the selected section, e.g. for volume adjustment or for the creation of simple fade-in and fade-out effects. Use Envelop to do more advanced amplitude shaping.

Right-click **Amplify** to open its settings dialog:

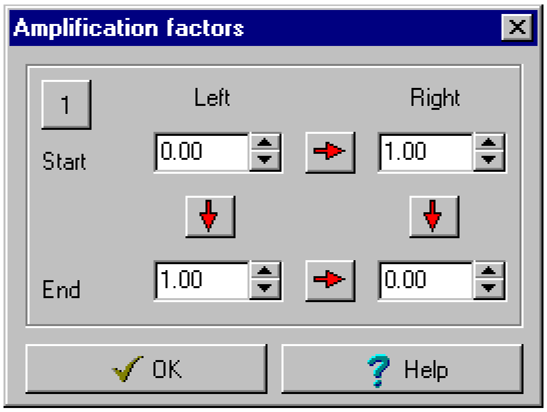

The edit boxes let you set **Start** and **End** amplification factors individually for each channel. The **Start** value is the amplification factor applied to the first sample in the selected section. The **End** value is the amplification factor applied to the last sample in the selected section. Intermediate values are computed by linear interpolation. Negative values (causing phase inversion) are allowed.

The picture shows the settings for a fade-in on the **Left** channel accompanied by a fade-out on the **Right** channel.

The **Right** column is ignored for monophonic files.

Use the arrow buttons to copy amplification factors to neighbouring boxes.

Clicking **1** causes all levels to be reset to 1.00, i.e. no amplification.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Amplify** will be executed.

Click 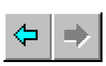 to abort and revert to the old settings.

Envelop

Envelop is used to impose an amplitude envelope on the selected section. Use Amplify for simple volume adjustments, fade-in and fade-out effects.

Right-click **Envelop** to open its settings dialog:

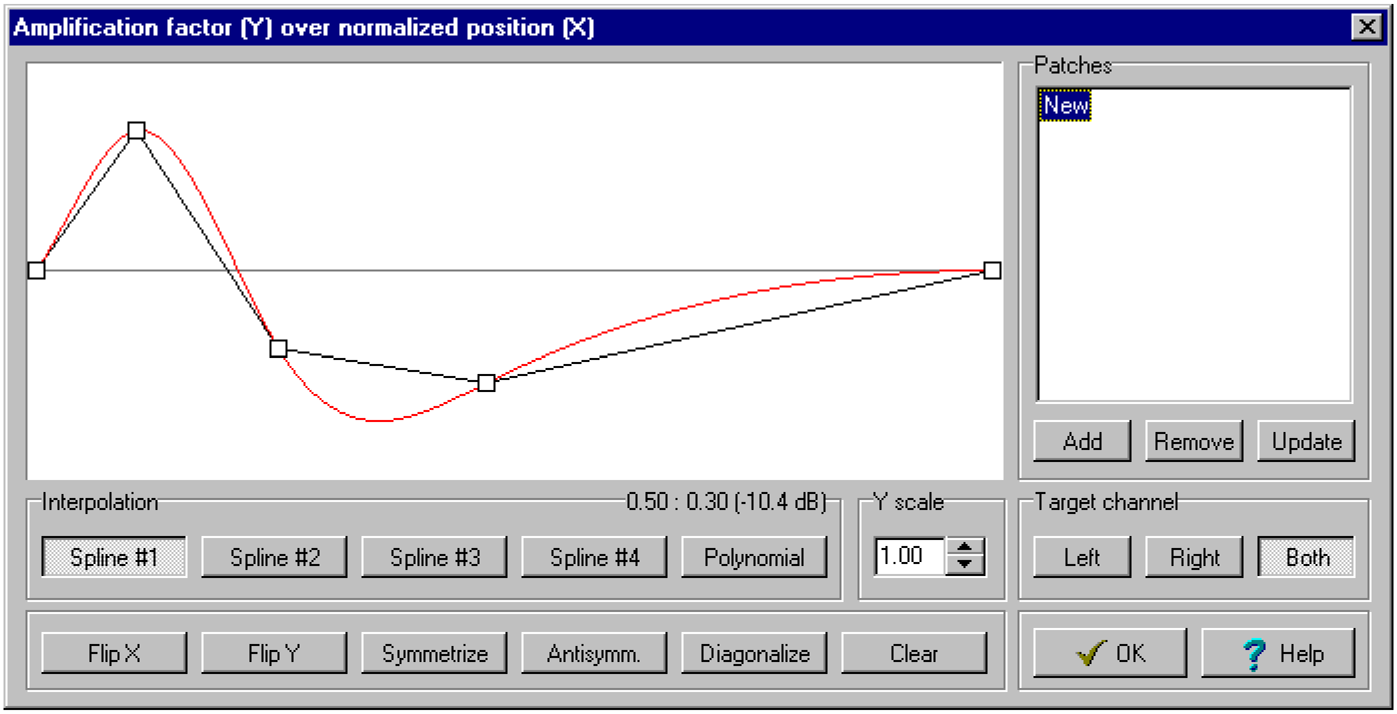

The envelope shape is entered and edited in the graph. You also have the editor buttons in the bottom box at your disposal.

Use the **Y scale** box to set the overall amplification scale.

By default, envelope values between nodes are computed by linear interpolation. If a button in the **Interpolation** box is depressed, two envelope curves will be displayed: the usual, linearly interpolated one (in black) and the selected, non-linear one actually imposed on the data (in red). You can choose between five different types of non-linear interpolation.

The cursor position display in the upper right corner of the **Interpolation** box is read as "normalized position : amplification factor" (time positions are normalized, i.e. 0.00 denotes the start and 1.00 the end of the selected section). Note that you can use negative amplification factors (phase inversion).

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own envelope Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Envelop** will be executed.

Click 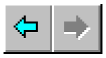 to abort and revert to the old settings.

Click any free spot in the graph to insert a new node there. Graph nodes can not overlap.

Place the mouse cursor on any node to highlight it:

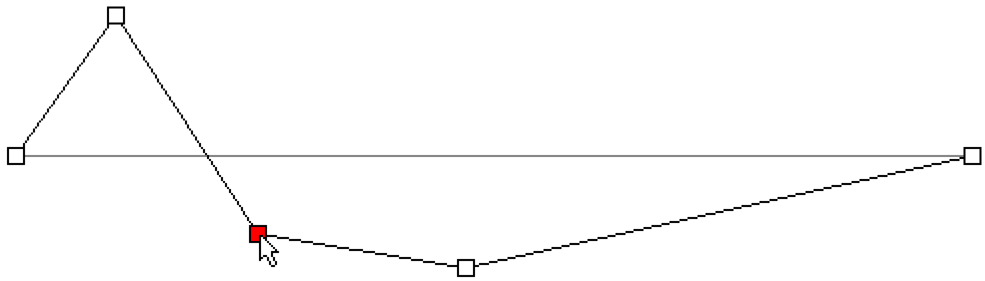

When a node is hightlighted it can be

- dragged to a different position (keep the left mouse button pressed while moving the mouse cursor around) and
- deleted: right-click to open the **Delete node** pop-up menu

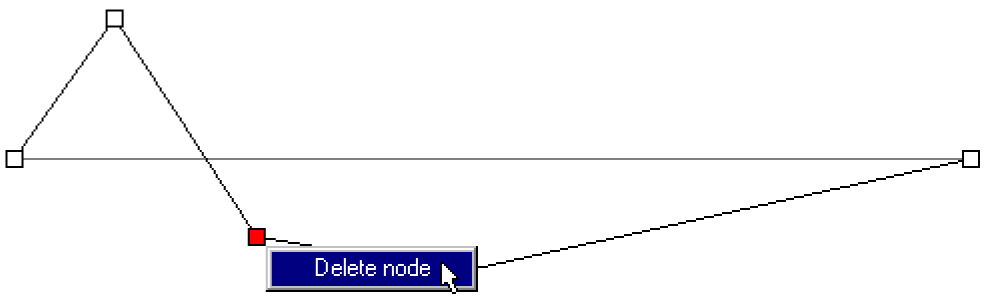

and click it to confirm deletion. Click anywhere outside the graph to close the pop-up menu and abort deletion.

The leftmost and rightmost nodes in the graph can neither be moved horizontally nor deleted.

Flip X Reverses (time-inverts) the graph.

Flip Y Inverts (changes the sign of) the graph.

Symmetrize Makes the graph time-symmetric (a time-symmetric graph is unaffected by reversal).

Antisymm. Makes the graph time-antisymmetric (reversing a time-antisymmetric graph is equivalent to inverting it).

Diagonalize Lines up all nodes from the graph's lower left corner to its upper right corner.

Clear Deletes all but the leftmost and rightmost nodes and moves those back to the horizontal axis.

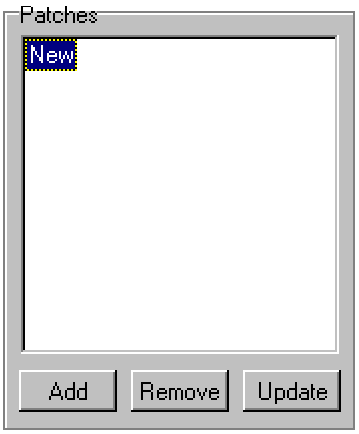

Click **Add** to save the current settings to a new patch.

New patches are always called "New". Patch names can be changed just like file names in the standard Windows file dialogs: select an entry in the list, wait a moment and click again to put it in edit mode: 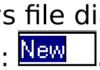

. Enter the new name, then hit Enter or click anywhere else in the list to exit edit mode. If you change your mind and want the old name back, you can abort editing by pressing Esc.

Select an entry in the list and click **Update** to write the current settings to an existing patch.

Use **Remove** to delete the selected patch.

Double-click any patch to load it. When a patch is selected, double-clicking anywhere in the list will also cause the patch to be loaded.

Acid WAV supports interpolation with cubic splines. A cubic spline is built up of third degree polynomials in such a way that the resulting function and its first two derivatives are continuous everywhere.

Any given set of points can be approximated by an infinite number of different cubic splines. In order to pick out a particular solution, two additional conditions must be specified. These are usually either

- "natural" boundary conditions: the spline must be flat outside of the interpolation interval

or

- "correct" boundary conditions: at the ends of the interpolation interval, the spline's slope must coincide with that of the function being interpolated (in practice computed by linear interpolation).

Acid WAV lets you combine the above conditions as follows:

- **Spline #1**: "correct" boundary conditions at both ends.
- **Spline #2**: "natural" boundary condition at the left end, "correct" boundary condition at the right end.
- **Spline #3**: "correct" boundary condition at the left end, "natural" boundary condition at the right end.
- **Spline #4**: "natural" boundary conditions at both ends.

You can also choose to interpolate all points with a single, higher order **Polynomial**. Polynomial interpolation tends to suffer from wider oscillations than spline interpolation, but can yield better results in special cases.

Maximize

Maximize is used to amplify the selected section so as to exploit as much as possible of the available dynamic range (-32768 to 32767 for 16 bit files, 0 to 255 for 8 bit files).

If the Statistics window indicates the presence of a DC offset in the sound data, you can make more room for maximization using the Offset function first.

Right-click **Maximize** to open its settings dialog:

Clicking **Left** or **Right** causes only one channel to be maximized, independently of the other channel. If **Both** channels are maximized, the relation between their levels is preserved (usually meaning that only one of them will be fully maximized).

The **Target channel** setting is ignored when maximizing monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Maximize** will be executed.

Click 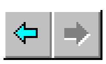 to abort.

Offset

Offset is used to shift all sample values in the selected section by a constant.

Right-click **Offset** to open its settings dialog:

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

The **Operation** box lets you choose between

- removing any DC component (**Center average**) present in the sound

- making equal room between the largest sample value and the "ceiling" (32767 for 16 bit files, 255 for 8 bit files) on one hand and the smallest sample value and the "floor" (-32768 for 16 bit files, 0 for 8 bit files) on the other (**Center extrema**) - this optimizes the data for a Maximize, but note that it will also introduce a DC component which might become troublesome later on

- specifying an arbitrary offset (**Manual offset** - the button must be depressed for the edit box to be enabled).

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Offset** will be executed.

Click 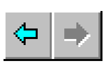 to abort and revert to the old settings.

Invert

Invert is used to flip the phase of the selected section. It's equivalent to (but more convenient than) Amplifying with all levels set to -1.00.

Common applications (together with Copy and Paste mix) include removing a known sound component from a mix and extracting differences between sounds.

Right-click **Invert** to open its settings dialog:

Clicking **Left** or **Right** causes only one channel to be inverted.

The **Target channel** setting is ignored when inverting monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Invert** will be executed.

Click 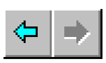 to abort.

Gate

Gate is used to silence those parts of the selected section where the amplitude falls below a given threshold. Use Compress / Expand to implement more articulate responses to amplitude levels.

A common application is the suppression of background noise in recordings containing (almost) silent sections. Another one is in effects, e.g. gated reverb.

Right-click **Gate** to open its settings dialog:

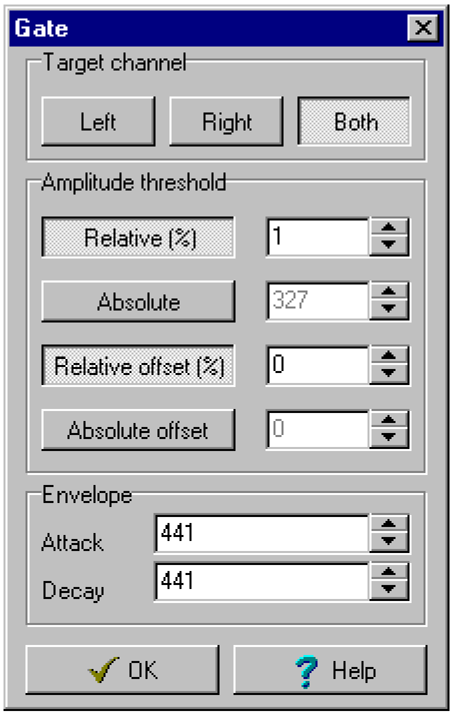

Selecting **Left** or **Right** causes only one channel to be gated. The **Target channel** setting is ignored for monophonic files.

The **Amplitude threshold** box lets you specify the amplitude threshold either in **Relative** terms (as a percentage of the full amplitude range) or as an **Absolute** value. You can also offset the threshold, e.g. to compensate for a known DC component (see Statistics). Keep in mind that the max amplitude of an 8 bit file is 128, that of a 16 bit file 32768.

Abrupt clipping when the amplitude crosses the threshold level is avoided by specifying non-zero **Attack** (fade in) and **Decay** (fade out) times in the **Envelope** box.

When the dialog is called up from the main window, durations are expressed in the main window's time unit.

When the dialog is called up from the script editor, durations are expressed in milliseconds.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Gate** will be executed.

Click 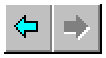 to abort and revert to the old settings.

Clip

Clip is used to limit sample values in the selected section.

Right-click **Clip** to open its settings dialog:

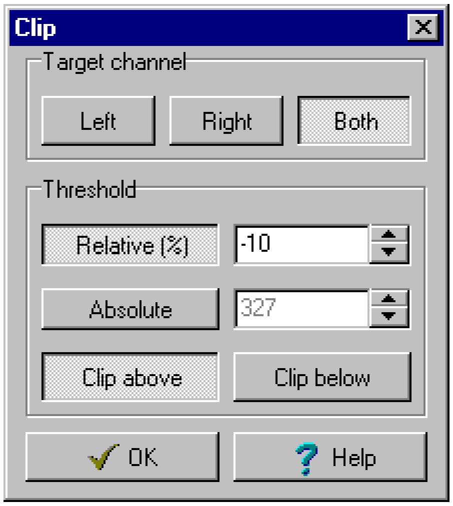

Selecting **Left** or **Right** causes only one channel to be clipped. The **Target channel** setting is ignored for monophonic files.

The **Threshold** box lets you specify the amplitude threshold either in **Relative** terms (as a percentage of the full amplitude range) or as an **Absolute** value.

Note that the threshold is signed, i.e. it's a sample value, not an amplitude. In order to clip on amplitude, you need to apply **Clip** twice. You could for instance **Clip above** a 10% threshold and then **Clip below** a -10% threshold, obtaining a 10% amplitude threshold clipping.

Keep in mind that the max amplitude of an 8 bit file is 128, that of a 16 bit file 32768.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Clip** will be executed.

Click 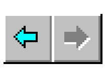 to abort and revert to the old settings.

Distort

Distort runs each sample in the selected section through a user-defined function.

Right-click **Distort** to open its settings dialog:

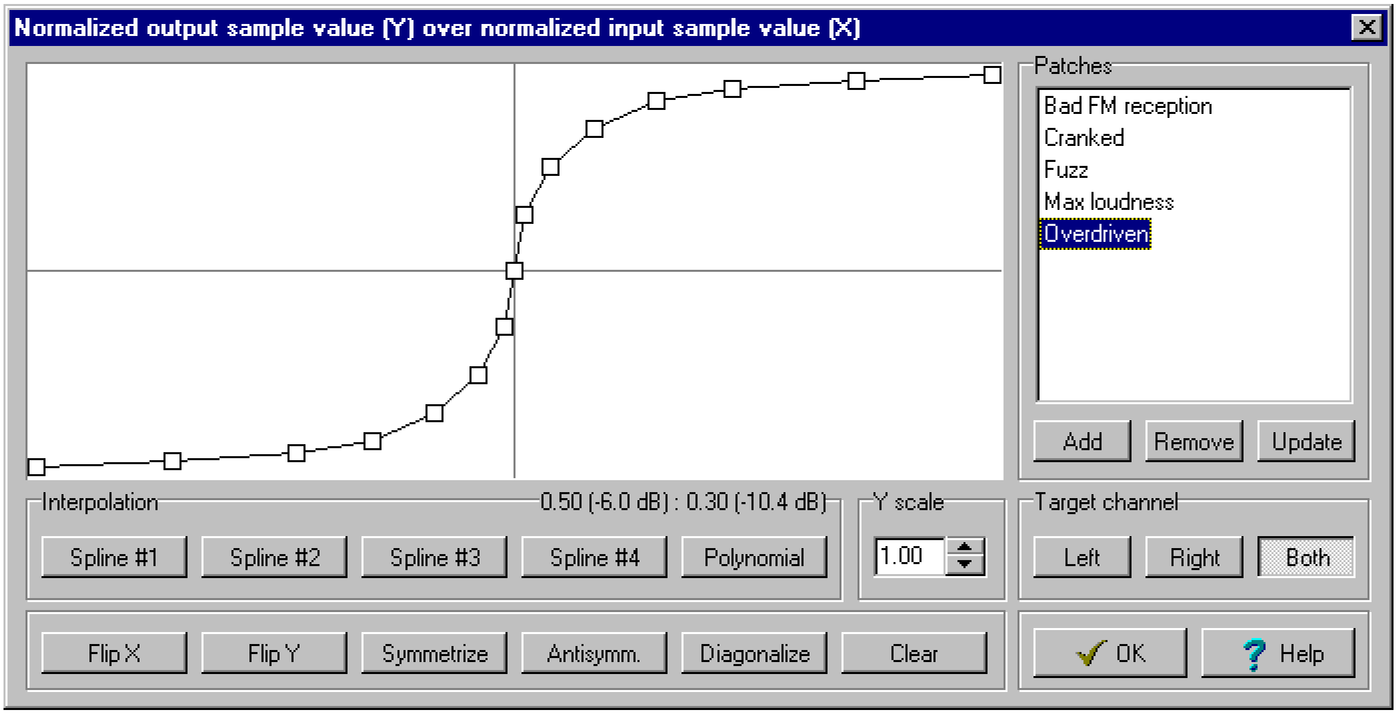

The distortion function shape is entered and edited in the graph. You also have the editor buttons in the bottom box at your disposal.

Note that the graph shows (normalized) sample values, not amplitudes. This allows you to create asymmetric distortion functions behaving differently above and below the silence level. See Compress / Expand for an amplitude-based alternative.

The cursor position display in the upper right corner of the **Interpolation** box is read as "normalized input sample value : normalized output sample value". Sample values are normalized to the range -1.0 to 1.0 in order to make distortion functions independent of the file format.

Use the **Y scale** box to set the overall amplification scale.

By default, function values between nodes are computed by linear interpolation. If a button in the **Interpolation** box is depressed, two curves will be displayed: the usual, linearly interpolated one (in black) and the selected, non-linear one actually applied to the data (in red). You can choose between five different types of non-linear interpolation.

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own distortion Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window,

Distort will be executed.

Click 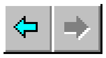 to abort and revert to the old settings.

Rectify

Rectify is used to invert samples with values below the silence level (0 for 16 bit files, 128 for 8 bit files) throughout the selected section.

Right-click **Rectify** to open its settings dialog:

Clicking **Left** or **Right** causes only one channel to be rectified. The **Target channel** setting is ignored when rectifying monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Rectify** will be executed.

Click 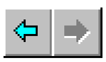 to abort.

Graphically,

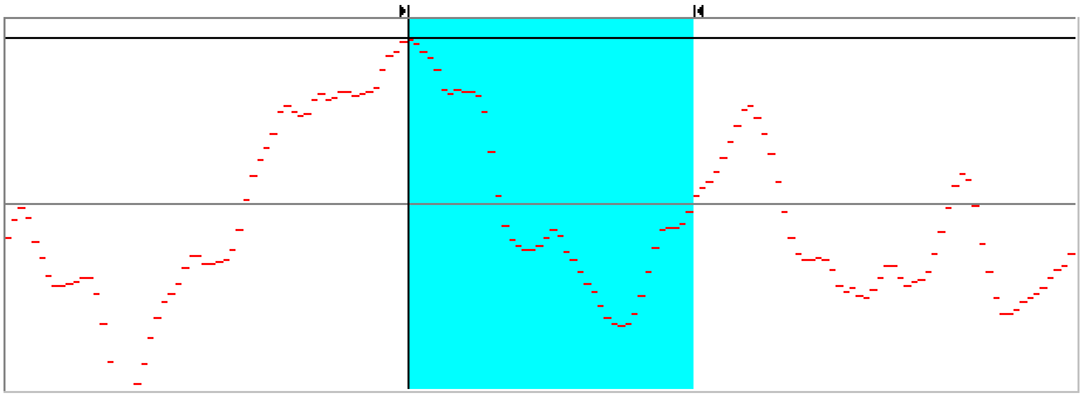

becomes

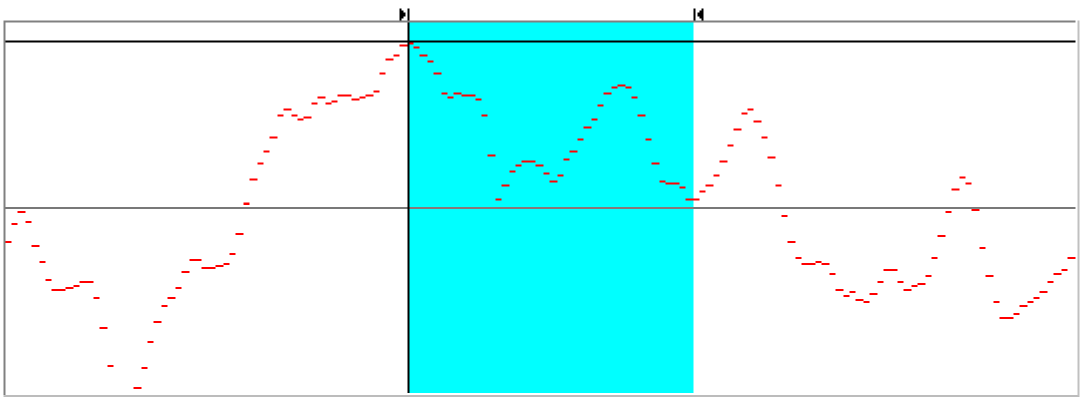

Quantize

Quantize is used to change the number of effective quantization levels in the selected section.

Right-click **Quantize** to open its settings dialog:

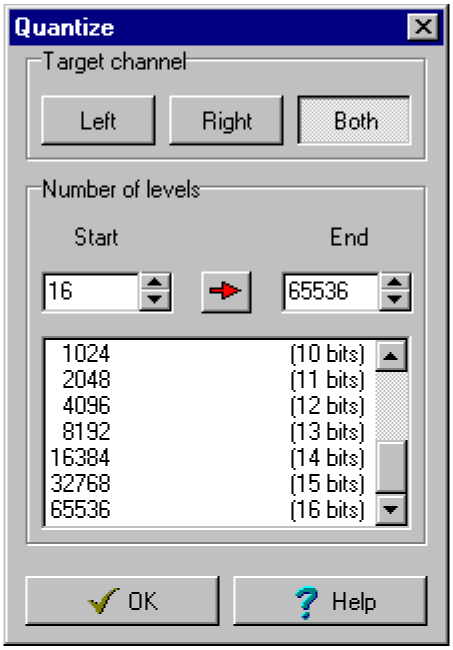

Selecting **Left** or **Right** causes only one channel to be quantized. The **Target channel** setting is ignored for monophonic files.

Set the number of effective quantization levels for the first sample in the **Start** edit box.

Set the number of effective quantization levels for the last sample in the **End** edit box.

The number of effective quantization levels for intermediate samples is computed by linear interpolation.

Click an entry in the list to make it the new **Start** value. Click the arrow to copy the **Start** value to the **End** value.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Quantize** will be executed.

Click 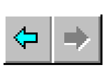 to abort and revert to the old settings.

The samples in digital audio files are **quantized**, i.e. they can not have arbitrary values.

8 bit samples have 256 possible values or **quantization levels** (all integers between 0 and 255).

16 bit samples have 65536 quantization levels (all integers between -32768 and 32767).

The Quantize function reduces the number of **effective quantization levels**, i.e. the number of quantization levels which are actually used.

For instance, reducing the number of effective quantization levels of a 16 bit file to 32768 causes every other actual quantization level to go unused. Sample values = 1 are changed to 0, sample values = 3 are changed to 2 and so on. (Trying to do the same to an 8 bit file has no effect: the number of effective quantization levels can never be larger than the number of actual quantization levels.)

Gated delete

Gated delete is used to remove those parts of the selected section which fail an amplitude threshold test.

A common application is the automated removal of leading and trailing silence from recordings.

Right-click **Gated delete** to open its settings dialog:

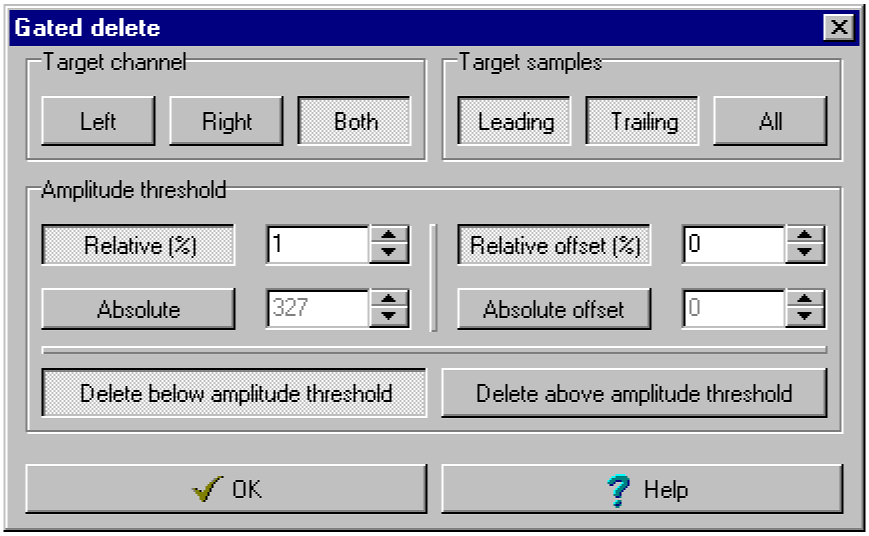

Selecting **Left** or **Right** causes only one channel to be tested and modified. The resulting difference in duration is compensated by padding the selected channel with silence, leaving the overall file length unchanged.

Selecting **Both** causes the test to be applied to the average value of both channels and samples from both channels to be removed together.

The **Target channel** setting is ignored for monophonic files.

The **Target samples** box lets you choose between the following modes of operation:

- **Leading.** Starting from the first sample in the selected section and moving forward, remove all samples that fail the amplitude threshold test and abort at the first sample that passes the test. Typically used to remove leading silence.
- **Trailing**. Starting from the last sample in the selected section and moving backward, remove all samples that fail the amplitude threshold test and abort at the first sample that passes the test. Typically used to remove trailing silence.
- **All**. Remove all samples that fail the amplitude threshold test. This can be used to create a "ripping" or "scratching" sound effect.

As shown in the picture, **Leading** and **Trailing** can be combined.

The **Amplitude threshold** box lets you specify the amplitude threshold either in **Relative** terms (as a percentage of the full amplitude range) or as an **Absolute** value. You can also offset the threshold, e.g. to compensate for a known DC component (see Statistics). Keep in mind that the max amplitude of an 8 bit file is 128, that of a 16 bit file 32768.

If **Delete below amplitude threshold** is depressed, the threshold test is the same used by the ordinary Gate function, i.e. low samples are deleted. If **Delete above amplitude threshold** is depressed, high samples are deleted.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Gated delete** will be executed.

Click 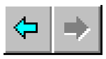 to abort and revert to the old settings.

Interpolate

Interpolate is used to replace all sample values in the selected section with a line connecting the start and end samples.

This can be a useful complement to deleting and filtering when removing pops, clicks and scratches from recordings.

Right-click **Interpolate** to open its settings dialog:

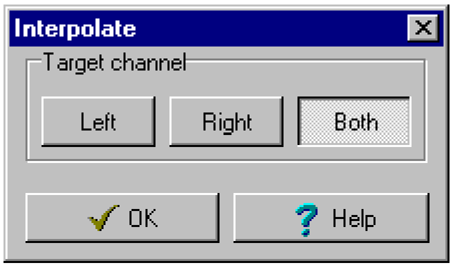

Clicking **Left** or **Right** causes only one channel to be interpolated.

The **Target channel** setting is ignored when interpolating monophonic files.

Click **OK** to confirm the current settings or any button in the **Target channel** box to update them. If the dialog was called up from the main window, **Interpolate** will be executed.

Click 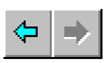 to abort.

Graphically,

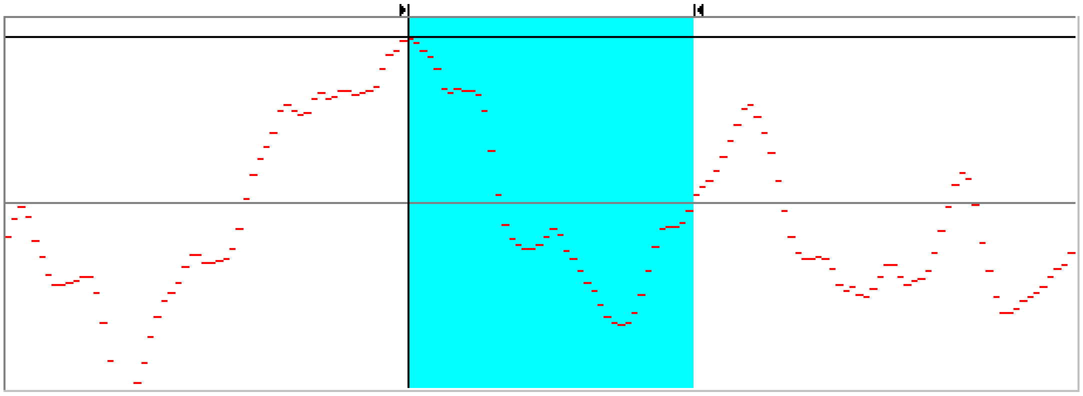

becomes

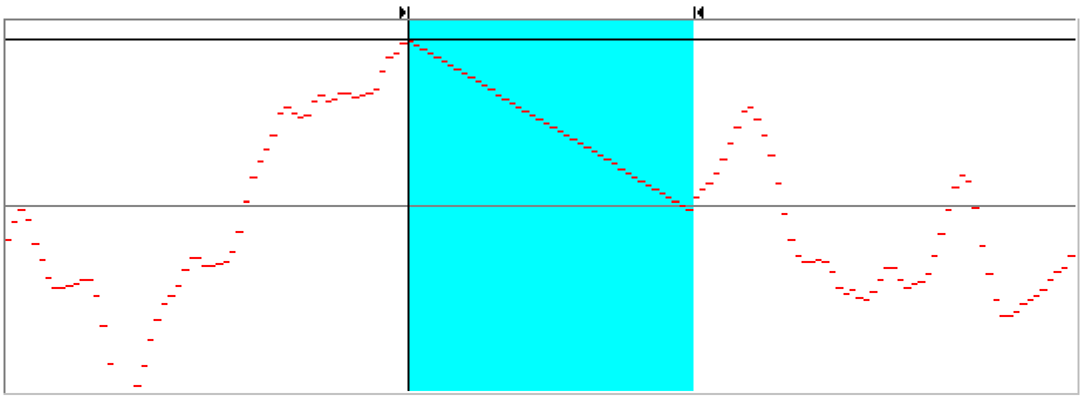

Compress / Expand

Compress / Expand modifies the moving average of amplitudes in the selected section according to a user-defined function. It can be used to compress, limit, expand and gate sounds.

Right-click **Compress / Expand** to open its settings dialog:

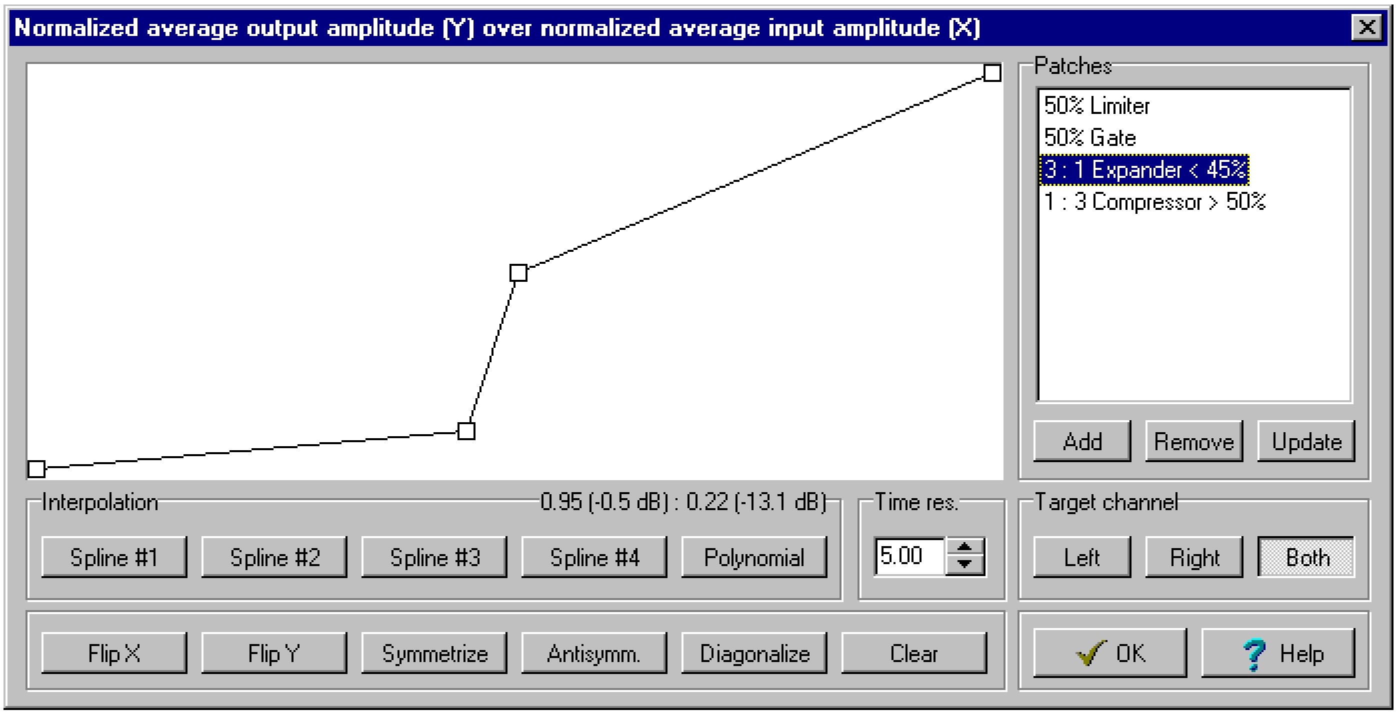

The amplitude function shape is entered and edited in the graph. You also have the editor buttons in the bottom box at your disposal.

The cursor position display in the upper right corner of the **Interpolation** box is read as "normalized average input amplitude : normalized average output amplitude". Amplitudes are normalized to the range -1.0 to 1.0 in order to make amplitude functions independent of the file format.

Use the **Time res.** box to set the time resolution, i.e. the length of the moving frame in which average amplitudes are computed. Reducing the time resolution to 1 sample turns **Compress / Expand** into a close relative of the Distort function (the difference being partly that **Distort** works on sample values rather than on amplitudes, partly that it doesn't treat stereo channels jointly - see below).

When the dialog is called up from the main window, the time resolution is expressed in the main window's time unit.

When the dialog is called up from the script editor, the time resolution is expressed in milliseconds.

By default, function values between nodes are computed by linear interpolation. If a button in the **Interpolation** box is depressed, two curves will be displayed: the usual, linearly interpolated one (in black) and the selected, non-linear one actually applied to the data (in red). You can choose between five different types of non-linear interpolation.

Selecting **Left** or **Right** causes only one channel to be modified, independently of the other. Select **Both** to average the amplitudes of left and right channels jointly, applying the same amplification factor to both. This prevents unnatural-sounding effects like large differences in background noise levels between channels.

The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own compress / expand Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Compress / Expand** will be executed.

Click 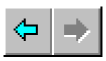 to abort and revert to the old settings.

Compressors reduce the dynamic range of signals. In their simplest form, they are implemented as amplifiers with two gain levels. Below a given amplitude threshold, the gain is unity (i.e. the sound is let through unchanged). Above the threshold, the gain is less than unity. The result is to boost low amplitude signals relative to high amplitude signals.

A **limiter** is an extreme compressor characterized by a maximal output level clamped to the threshold level.

Expanders increase the dynamic range of signals. Like compressors, simple expanders can be implemented as amplifiers with two gain levels. Above a given amplitude threshold, the gain is unity (i.e. the sound is let through unchanged). Below the threshold, the gain is less than unity. The result is to boost high amplitude signals relative to low amplitude signals.

A gate is an extreme expander characterized by total suppression below the threshold level.

Frequency functions

The **Frequency** editing group contains the following functions:

- Low pass
- High pass
- Band pass
- Band stop
- Resonate
- Notch
- Equalize
- Convolve
- Vocode
- Resample
- Modulate
- Pitch shift
- Reduce noise

Low pass

Low pass is used to remove frequencies above a given cutoff value from the selected section. It's the opposite of High pass.

See also Band pass, Band stop, Resonate, Notch, Equalize and Convolve.

Right-click **Low pass** to open its settings dialog:

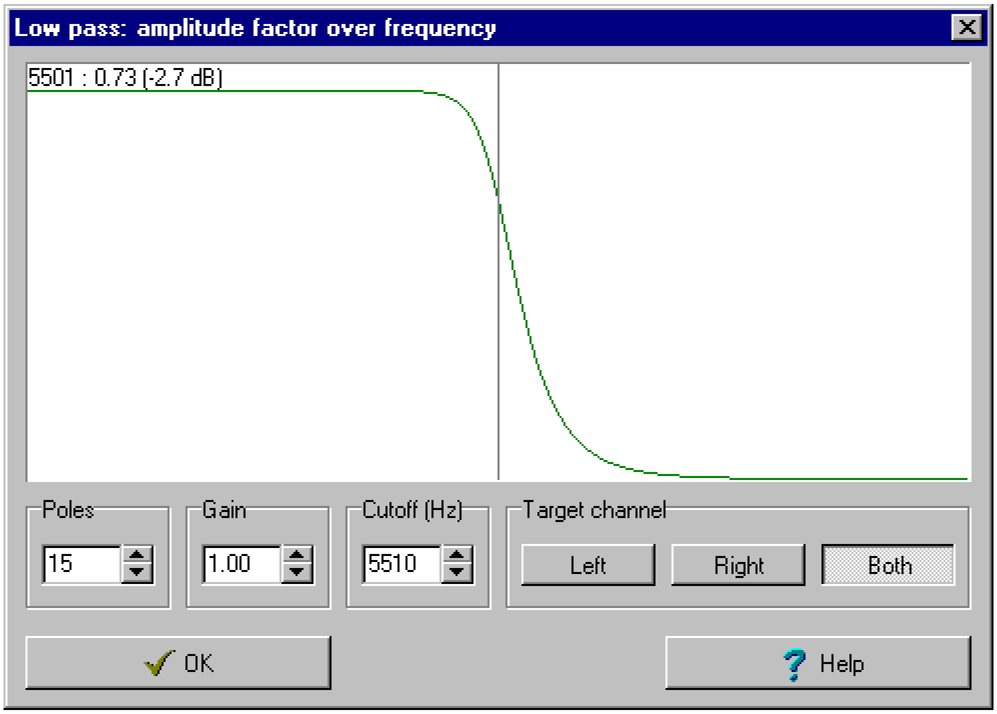

The graph shows the frequency response of the filter. The vertical line marks the cutoff frequency. The cursor position display in the graph's upper left corner is read as "frequency : amplification factor".

If the dialog is called up from Acid WAV's main window when a file is loaded, frequencies are expressed in Hertz (as in the picture). If the dialog is called up from the script editor or from the main window when no file is loaded, frequencies are normalized: the sampling rate is then 1.0 and the Nyqvist frequency (the highest reproducible frequency, and therefore the right end of the graph) is 0.5.

The cutoff frequency can be set by clicking the graph or directly in the **Cutoff** box.

The narrowness of the transition band (i.e. how fast the frequency response drops off from full to zero)   depends on the number of **Poles**. If you don't know what poles are, think of them as filter elements. Adding more elements to the filter makes it more precise. The improvement drops off for each new pole, and the finite precision of computer arithmetics sets an absolute limit at about 50 poles. In most cases, there is little point to going beyond 30 poles.

Use the **Gain** box to set an overall amplification factor for the filter.

Selecting **Left** or **Right** causes only one channel to be filtered. The **Target channel** setting is ignored when filtering monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Low pass** will be executed.

Click 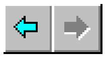 to abort and revert to the old settings.

High pass

High pass is used to remove frequencies below a given cutoff value from the selected section. It's the opposite of Low pass.

See also Band pass, Band stop, Resonate, Notch, Equalize and Convolve.

Right-click **High pass** to open its settings dialog:

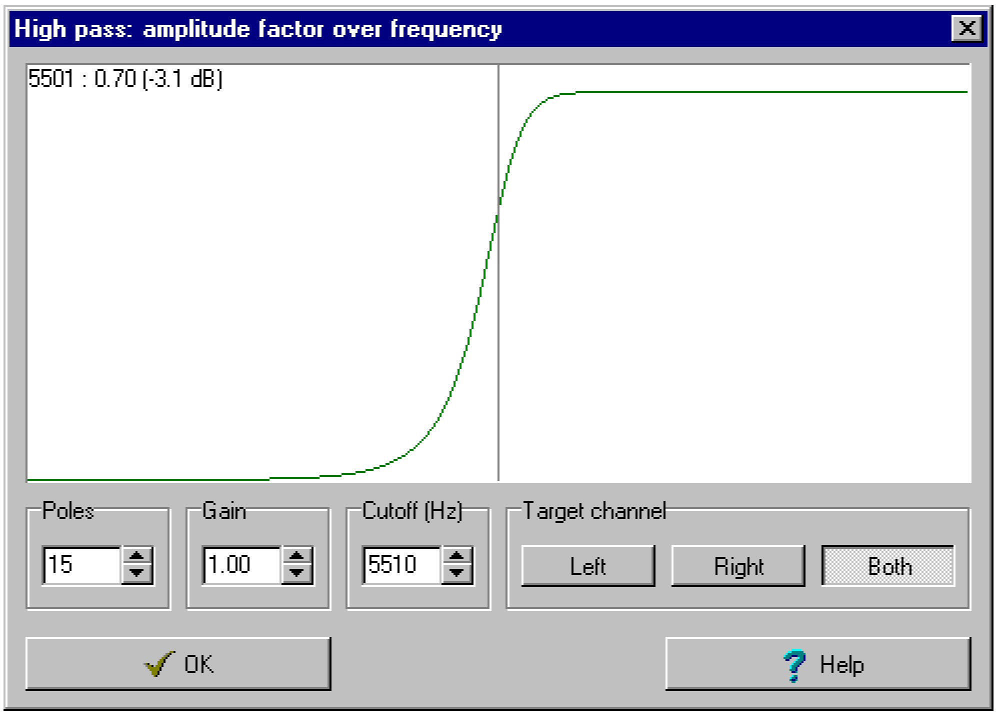

The graph shows the frequency response of the filter. The vertical line marks the cutoff frequency. The cursor position display in the graph's upper left corner is read as "frequency : amplification factor".

If the dialog is called up from Acid WAV's main window when a file is loaded, frequencies are expressed in Hertz (as in the picture). If the dialog is called up from the script editor or from the main window when no file is loaded, frequencies are normalized: the sampling rate is then 1.0 and the Nyqvist frequency (the highest reproducible frequency, and therefore the right end of the graph) is 0.5.

The cutoff frequency can be set by clicking the graph or directly in the **Cutoff** box.

The narrowness of the transition band (i.e. how fast the frequency response grows from zero to full) depends on the number of **Poles**. If you don't know what poles are, think of them as filter elements. Adding more elements to the filter makes it more precise. The improvement drops off for each new pole, and the finite precision of computer arithmetics sets an absolute limit at about 50 poles. In most cases, there is little point to going beyond 30 poles.

Use the **Gain** box to set an overall amplification factor for the filter.

Selecting **Left** or **Right** causes only one channel to be filtered. The **Target channel** setting is ignored when filtering monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **High pass** will be executed.

Click 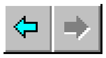 to abort and revert to the old settings.

Band pass

Band pass is used to remove frequencies outside a given frequency band from the selected section. It's the opposite of Band stop.

See also High pass, Low pass, Resonate, Notch, Equalize and Convolve.

Right-click **Band pass** to open its settings dialog:

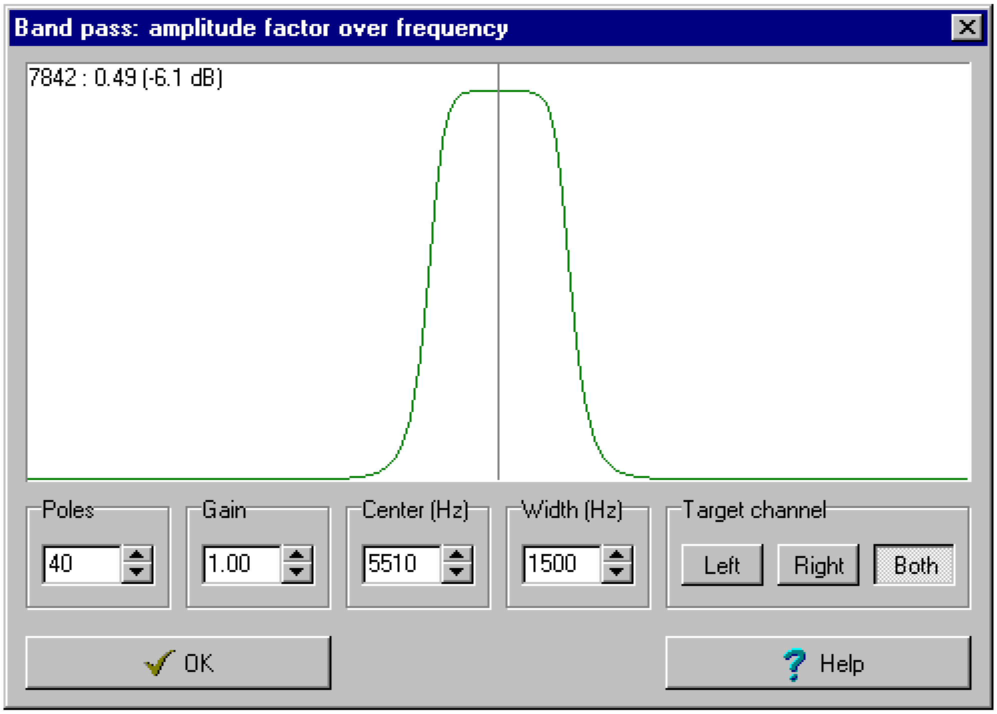

The graph shows the frequency response of the filter. The vertical line marks the center of the passband. The cursor position display in the graph's upper left corner is read as "frequency : amplification factor".

If the dialog is called up from Acid WAV's main window when a file is loaded, frequencies are expressed in Hertz (as in the picture). If the dialog is called up from the script editor or from the main window when no file is loaded, frequencies are normalized: the sampling rate is then 1.0 and the Nyqvist frequency (the highest reproducible frequency, and therefore the right end of the graph) is 0.5.

The center of the passband can be set by clicking the graph or directly in the **Center** box. The width is set in the **Width** box.

The narrowness of the transition bands (i.e. how fast the frequency response goes from zero to full at the low end and from full to zero at the high end of the passband)   depends on the number of **Poles**. If you don't know what poles are, think of them as filter elements. Adding more elements to the filter makes it more precise. The improvement drops off for each new pole, and the finite precision of computer arithmetics sets an absolute limit at about 50 poles. In most cases, there is little point to going beyond 40 poles.

Use the **Gain** box to set an overall amplification factor for the filter.

Selecting **Left** or **Right** causes only one channel to be filtered. The **Target channel** setting is ignored when filtering monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Band pass** will be executed.

Click 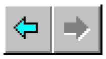 to abort and revert to the old settings.

Band stop

Band stop is used to remove frequencies inside a given frequency band from the selected section. It's the opposite of Band pass.

Use Notch to remove a narrow set of frequencies. See also High pass, Low pass, Resonate, Equalize and Convolve.

Right-click **Band stop** to open its settings dialog:

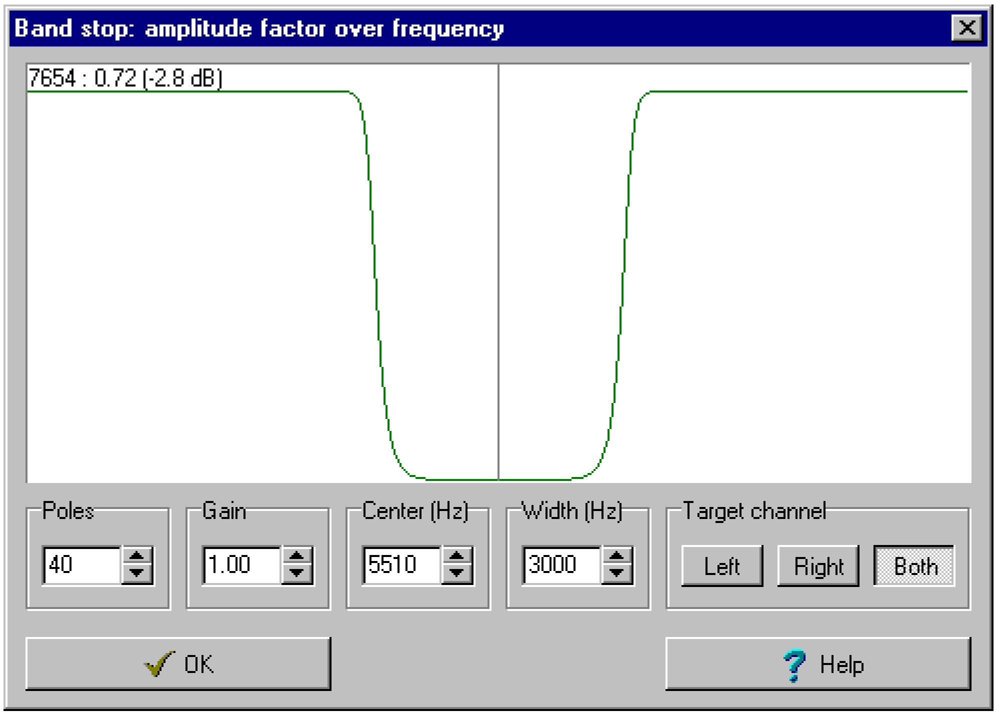

The graph shows the frequency response of the filter. The vertical line marks the center of the stopband. The cursor position display in the graph's upper left corner is read as "frequency : amplification factor".

If the dialog is called up from Acid WAV's main window when a file is loaded, frequencies are expressed in Hertz (as in the picture). If the dialog is called up from the script editor or from the main window when no file is loaded, frequencies are normalized: the sampling rate is then 1.0 and the Nyqvist frequency (the highest reproducible frequency, and therefore the right end of the graph) is 0.5.

The center of the stopband can be set by clicking the graph or directly in the **Center** box. The width is set in the **Width** box.

The narrowness of the transition bands (i.e. how fast the frequency response goes from full to zero at the low end and from zero to full at the high end of the stopband) depends on the number of **Poles**. If you don't know what poles are, think of them as filter elements. Adding more elements to the filter makes it more precise. The improvement drops off for each new pole, and the finite precision of computer arithmetics sets an absolute limit at about 50 poles. In most cases, there is little point to going beyond 40 poles.

Use the **Gain** box to set an overall amplification factor for the filter.

Selecting **Left** or **Right** causes only one channel to be filtered. The **Target channel** setting is ignored when filtering monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Band stop** will be executed.

Click 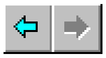 to abort and revert to the old settings.

Resonate

Resonate is used to sweep a lowpass filter over the selected section while emphasizing frequencies around the moving cutoff value.

Resonator sweeps are among the most common "analog" synth effects.

See also High pass, Low pass, Band pass, Band stop, Equalize and Convolve.

Right-click **Resonate** to open its settings dialog:

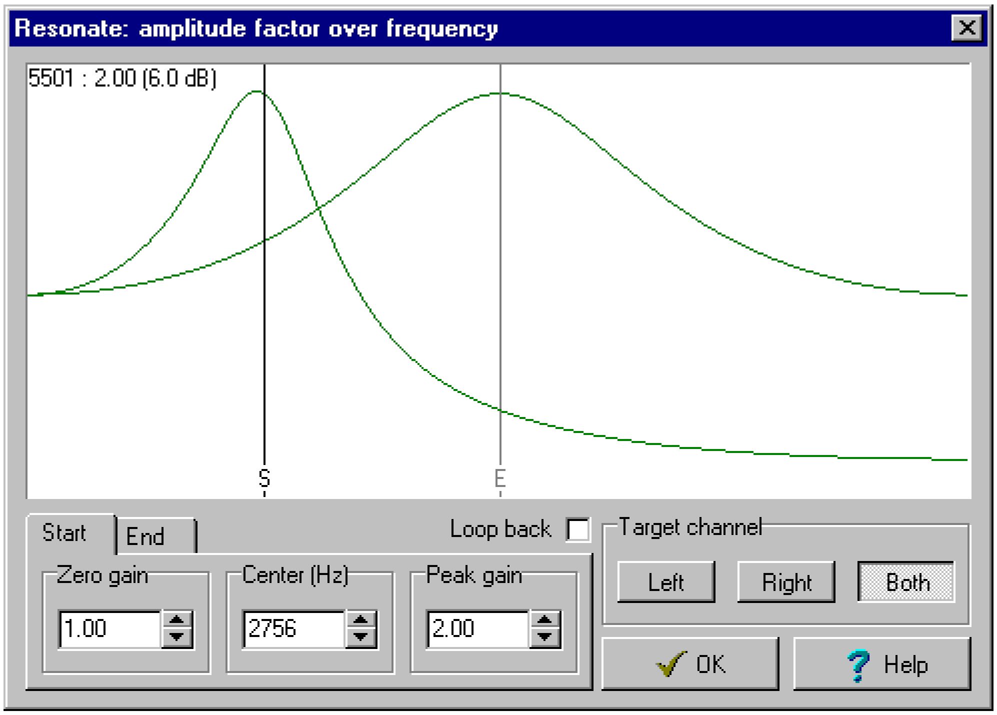

The graph shows the frequency response of the filter for the **Start** and **End** settings. The vertical lines mark the cutoff frequencies (labeled **S** for Start, **E** for End). The cutoff frequency for the selected tab is highlighted.

If **Loop back** is unchecked, the filter will be swept from the **Start** settings (at the first sample in the selected file section) to the **End** settings (at the last sample in the selected file section). All settings are linearly interpolated from **Start** to **End**. If **Loop back** is checked, the filter will reach the **End** settings halfway through the selected file section and then evolve back toward the **Start** settings.

For the selected tab,

- the center frequency can be set by clicking the graph or directly in the **Center** box,
- the amplification factor at the cutoff frequency is set in the **Peak gain** box, and
- the amplification factor at 0 Hz is set in the **Zero gain** box.

The cursor position display in the graph's upper left corner is read as "frequency : amplification factor".

If the dialog is called up from Acid WAV's main window when a file is loaded, frequencies are expressed in Hertz (as in the picture). If the dialog is called up from the script editor or from the main window when no file is loaded, frequencies are normalized: the sampling rate is then 1.0 and the Nyqvist frequency (the highest reproducible frequency, and therefore the right end of the graph) is 0.5.

Selecting **Left** or **Right** causes only one channel to be filtered. The **Target channel** setting is ignored when filtering monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Resonate** will be executed.

Click 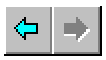 to abort and revert to the old settings.

Notch

Notch is used to remove a narrow set of frequencies from the selected section.

Use Band stop to remove a frequency band. See also High pass, Low pass, Band pass, Resonate, Equalize and Convolve.

Right-click **Notch** to open its settings dialog:

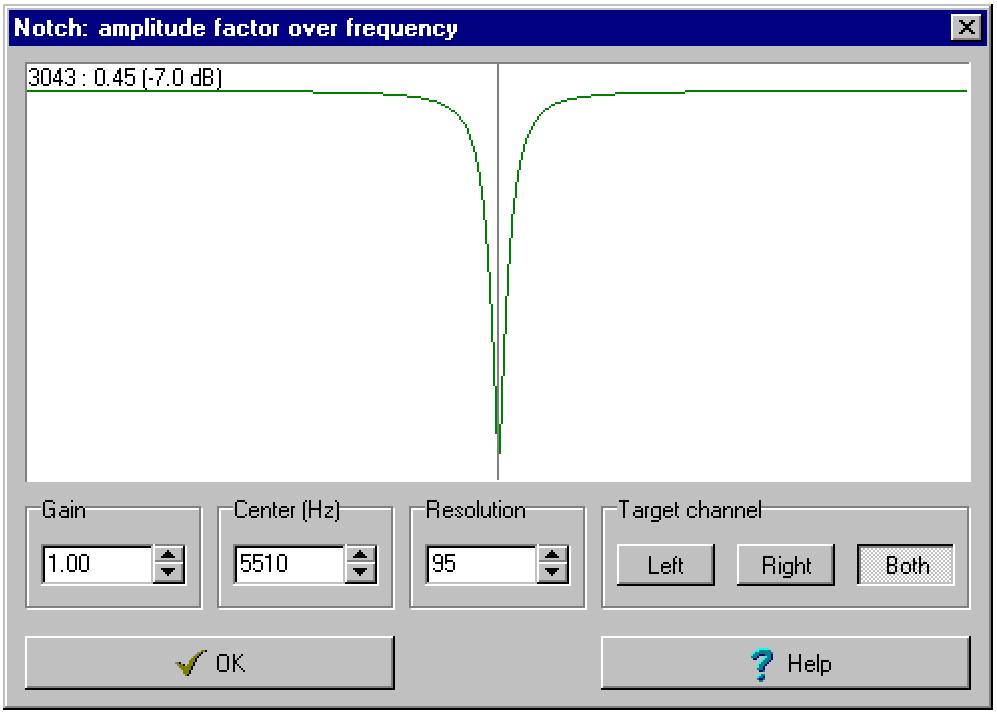

The graph shows the frequency response of the filter. The vertical line marks the center of the frequency range to remove. The cursor position display in the graph's upper left corner is read as "frequency : amplification factor".

If the dialog is called up from Acid WAV's main window when a file is loaded, frequencies are expressed in Hertz (as in the picture). If the dialog is called up from the script editor or from the main window when no file is loaded, frequencies are normalized: the sampling rate is then 1.0 and the Nyqvist frequency (the highest reproducible frequency, and therefore the right end of the graph) is 0.5.

The center frequency can be set by clicking the graph or directly in the **Center** box.

The width of the frequency range to be removed is controlled with the **Resolution** box. The highest resolution setting (corresponding to minimum width) is 99.

Use the **Gain** box to set an overall amplification factor for the filter.

Selecting **Left** or **Right** causes only one channel to be filtered. The **Target channel** setting is ignored when filtering monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Notch** will be executed.

Click 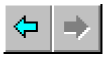 to abort and revert to the old settings.

Equalize

Equalize is used to modify the amplitudes of up to 10 x 2 independent frequency bands in the selected section.

See also High pass, Low pass, Band pass, Band stop, Resonate, and Notch.

Right-click **Equalize** to open its settings dialog:

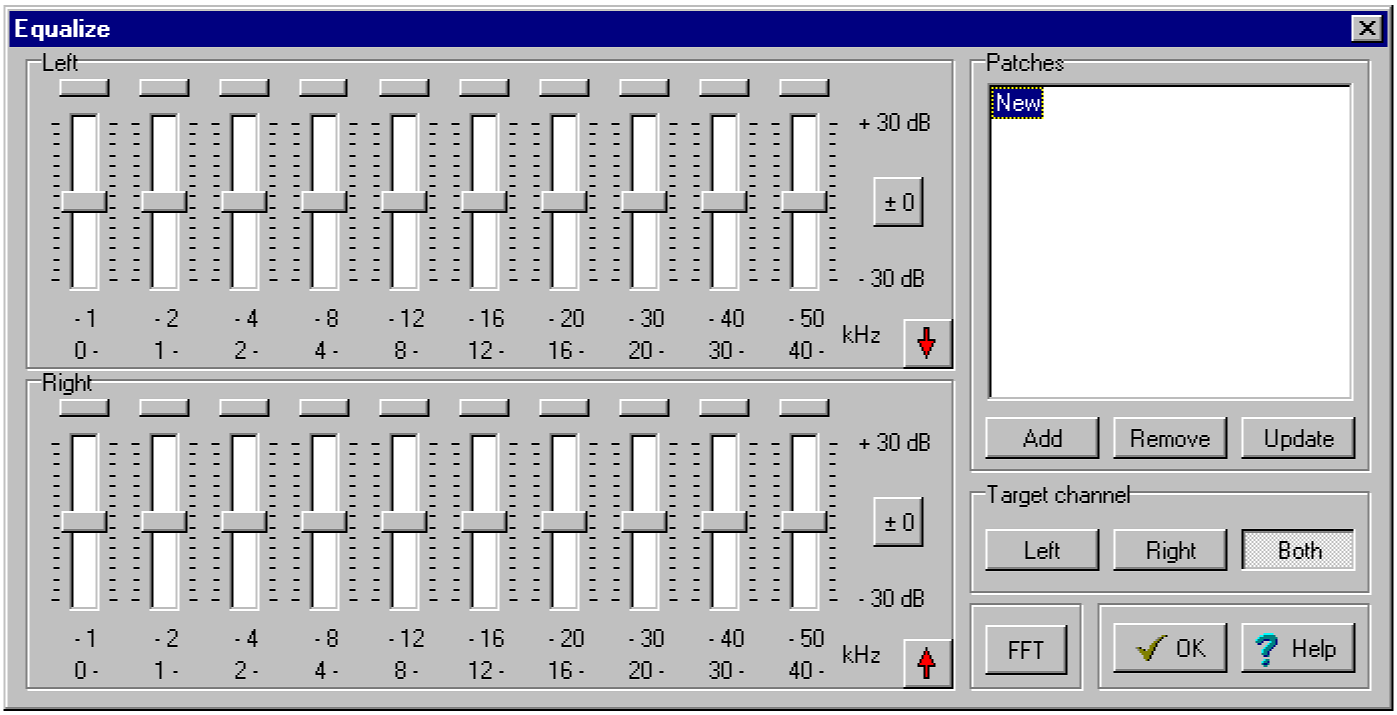

Use the gauges to set the desired amplification. The button above each gauge resets the gauge to 0 dB (no volume change).

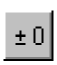

Resets all bands in the channel to 0 dB.

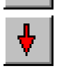
Copies the left channel settings to the right channel.

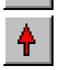
Copies the right channel settings to the left channel.

is used to select equalization algorithm:

- When **FFT** is depressed, the sound is decomposed using a Fast Fourier Transform. The main disadvantage of this option is that the FFT algorithm can introduce ripples at the ends of sections which start or end abruptly.

- When **FFT** is not depressed, the sound is decomposed using a bank of bandpass filters (low- and highpass filters at the low and high frequency ends). This is how hardware equalizers work. The main disadvantage of this option is that it can introduce larger phase delays than the FFT algorithm.

Selecting **Left** or **Right** causes only one channel to be equalized. The **Target channel** setting is ignored when equalizing monophonic files.

You can save all settings to (and restore them from) your own equalizer Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Equalize** will be executed.

Click 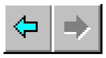 to abort and revert to the old settings.

Convolve

Convolve is used to run the selected section through a FIR (Finite Impulse Response) filter approximating a user-defined frequency response.

See also High pass, Low pass, Band pass, Band stop, Resonate, and Equalize.

Right-click **Convolve** to open its settings dialog:

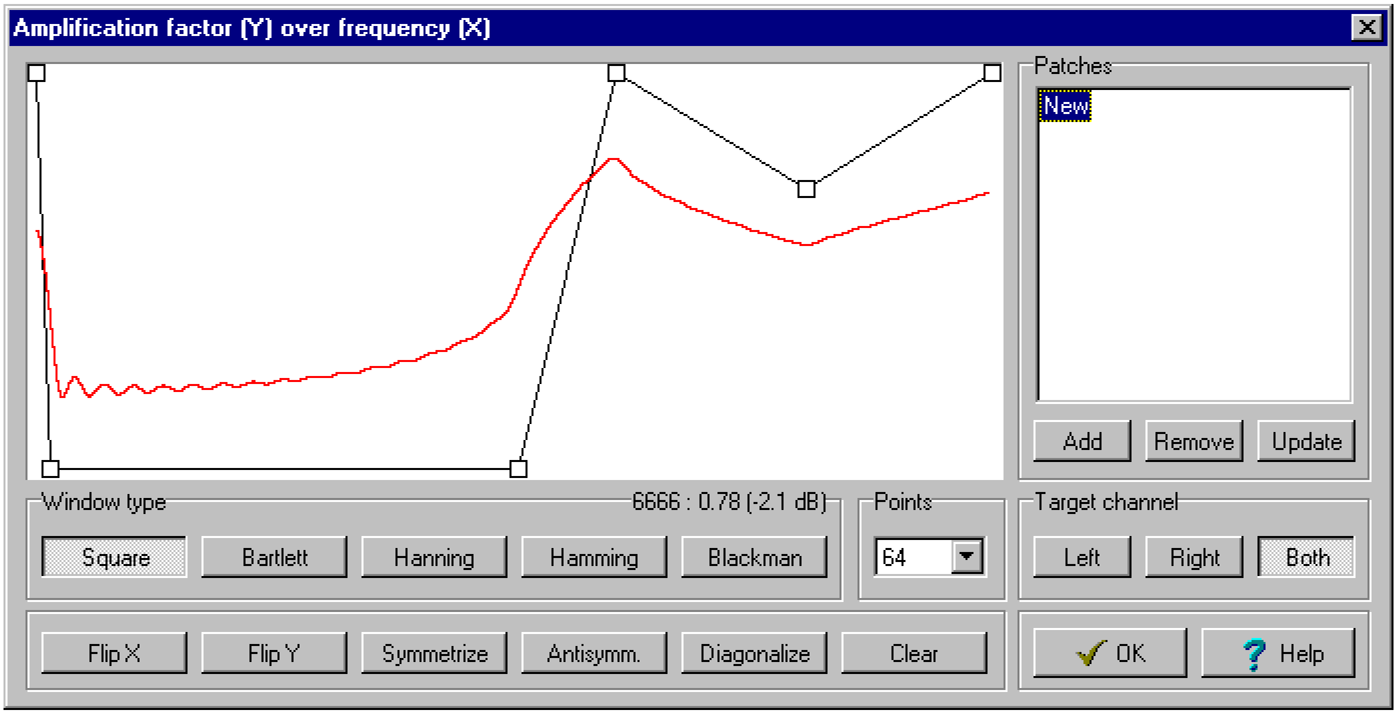

The desired frequency response is entered and edited in the graph. You also have the editor buttons in the bottom box at your disposal.

The actual frequency response is displayed in red. Apart from editing the desired response, you can affect it in two ways:

- Choose a different **Window type** to change the tradeoff between resolution (how closely the response curve follows sharp breaks in the desired curve) and smoothness. In the picture, a **Square** window is being used: this gives the highest resolution, but also the largest amount of rippling, as illustrated at the left end of the graph.
- Incrase the number of **Points**. This is the size of the FIR filter built by Acid WAV. A larger number of points allows a better (but also slower) filter to be created. The improvement drops off as more points are added, reflecting the fundamental limitations of FIR filters. The desired frequency response can usually not be reproduced exactly even with an inifinite number of points; in practice, even the max 2048 points allowed by **Convolve** are almost always overkill.

If the dialog is called up from Acid WAV's main window when a file is loaded, frequencies are expressed in Hertz (as in the picture). If the dialog is called up from the script editor or from the main window when no file is loaded, frequencies are normalized: the sampling rate is then 1.0 and the Nyquist frequency (the highest reproducible frequency, and therefore the right end of the graph) is 0.5.

Selecting **Left** or **Right** causes only one channel to be filtered. The **Target channel** setting is ignored when filtering monophonic files.

You can save all settings to (and restore them from) your own convolution Patches.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Convolve** will be executed.

Click 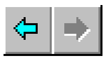 to abort and revert to the old settings.

Vocode

Vocode is used to impose the formant envelope of the clipboard (the modulator) on the selected section (the carrier).

The most common application is to let a human voice modulate a music instrument, making the instrument "speak" (listen to Kraftwerk's "Man Machine" for some classic examples). Many interesting effects can also be created with instruments serving both as carriers and as controllers. For instance, modulating a bass line with the accompanying kick and snare drums can add a lot of rhythmic life to a piece.

Right-click **Vocode** to open its settings dialog:

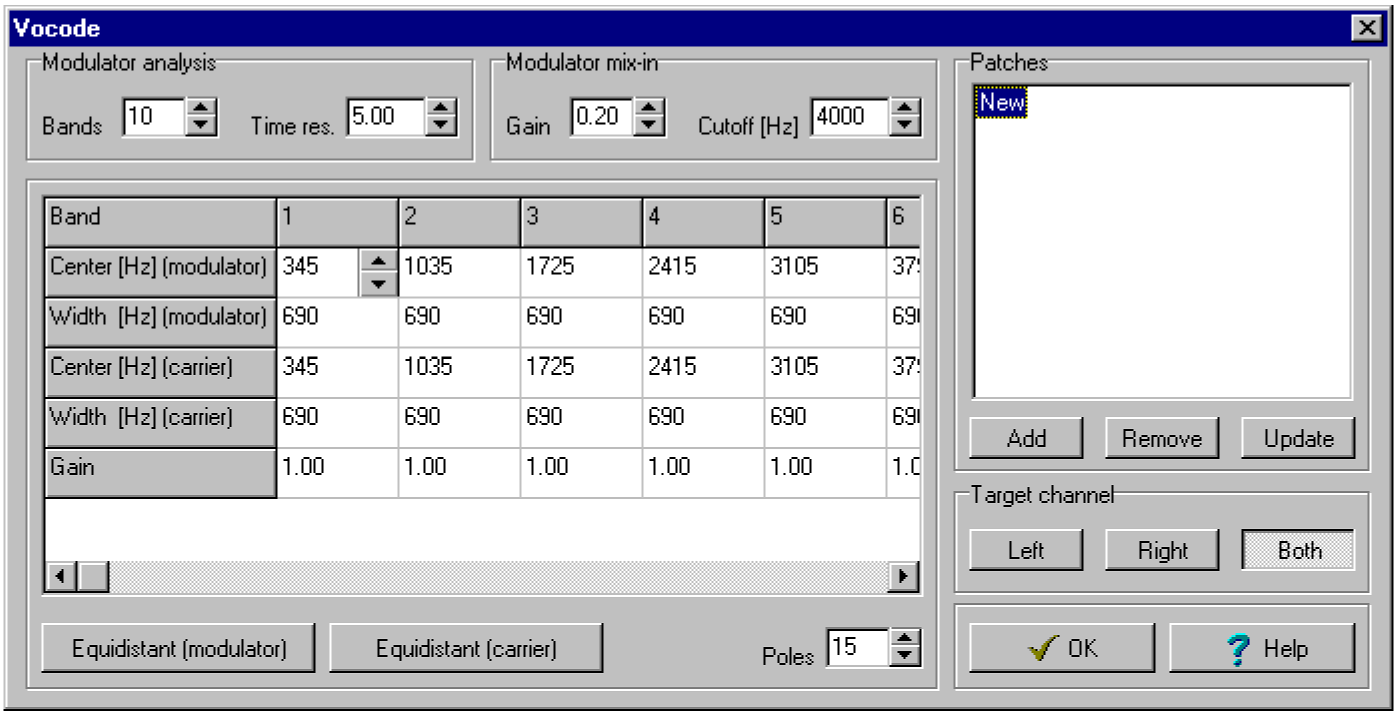

In order to understand this dialog, you need to understand how a vocoder works.

A vocoder is essentially a bank of bandpass filters controlling another bank of bandpass filters. The first bank slices up the modulator into frequency bands. The second filter bank slices up the carrier into frequency bands. Each modulator band is fed into a separate envelope follower, and the output of the envelope follower (the momentaneous amplitude of the controller band) is used to amplify (amplitude modulate) the corresponding carrier band.

The grid at the center of the **Vocode** dialog is where you specify how your modulator and carrier should be sliced. Each column describes a modulator band controlling a carrier band. **Center** and **Width** can be set independently for modulator and carrier, so there is nothing preventing you from matching up two completely different spectral regions.

IMPORTANT: The modulator is converted to the carrier's format before processing. Therefore, the carrier's sampling rate also determines which modulator filter bands are below the Nyqvist frequency (the highest reproducible frequency).

The **Gain** parameter allows you to amplify each carrier band by a constant factor. This is like having an equalizer cascaded after (and tuned to) the vocoder.

The number of modulator (and hence of carrier) **Bands** is set in the **Modulator analysis** box, where you also set the time resolution (**Time res.**) of the envelope followers. This value should not be too small (less than a few milliseconds): you want to modulate the carrier with an envelope, not with the amplitudes of individual samples.

In the grid display, you can see if a band column is enabled by looking at its header. If the column is disabled, its number is shown between parentheses, as in "(9)".

Enter a **Center** value for the modulator band in the rightmost enabled column and click **Equidistant (modulator)** to get a set of evenly spaced modulator bands covering the spectrum up to that value.

Enter a **Center** value for the carrier band in the rightmost enabled column and click **Equidistant (carrier)** to get a set of evenly spaced carrier bands covering the spectrum up to that value.

The **Poles** box is used to control the number of poles per bandpass filter, i.e. the frequency resolution of the filter banks. Generally speaking, a vocoder using narrow bands requires more poles than one using wide bands. This parameter (as well as **Center** and **Width**) works just like its counterpart in the Band pass dialog, so you can view the frequency response for each band there.

When the dialog is called up from the main window, the time resolution is expressed in the main window's time unit.

When the dialog is called up from the script editor, the time resolution is expressed in milliseconds.

Acid WAV's vocoder adds yet another twist to the basic vocoder concept by allowing the modulator to be highpass filtered and mixed into the final result. This can make speech modulating a music instrument more intelligible. You control this feature in the **Modulator mix-in** box. To shut it off, simply set the **Gain** to zero. The number of poles used for the bandpass filters does not affect the highpass filter.

Selecting **Left** or **Right** causes only one channel to be vocoded. The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own vocoder Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Vocode** will be executed.

Click 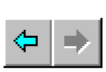 to abort and revert to the old settings.

Resample

Resample is used to change the frequency of the selected section by a constant or linearly varying factor. Use Modulate for more advanced frequency modulation.

Resampling modifies the duration of the sound by the inverse of the applied frequency factor. See Pitch shift for an alternative which leaves the sound's duration unchanged.

Right-click **Resample** to open its settings dialog:

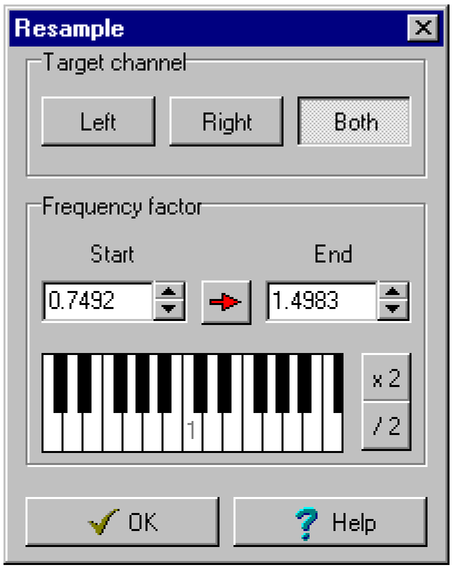

The frequency factor is linearly interpolated from the **Start** value (at the first sample in the selected section) to the **End** value (at the last sample in the selected section). The graphic keyboard can be used to enter **Start** factors corresponding to multiples of semitone shifts (useful e.g. when creating patches for samplers). Use the **x 2** and **/ 2** buttons to shift the **Start** factor by a full octave.

Click  to copy the **Start** factor to the **End** box.

Selecting **Left** or **Right** causes only one channel to be resampled. The resulting difference in duration is compensated by padding the shorter channel with silence, leaving the overall file length unchanged.

The **Target channel** setting is ignored for monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Resample** will be executed.

Click 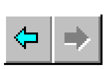 to abort and revert to the old settings.

Modulate

Modulate is used to impose a variable frequency factor on the selected section. Use Resample for constant or linear frequency changes.

Right-click **Modulate** to open its settings dialog:

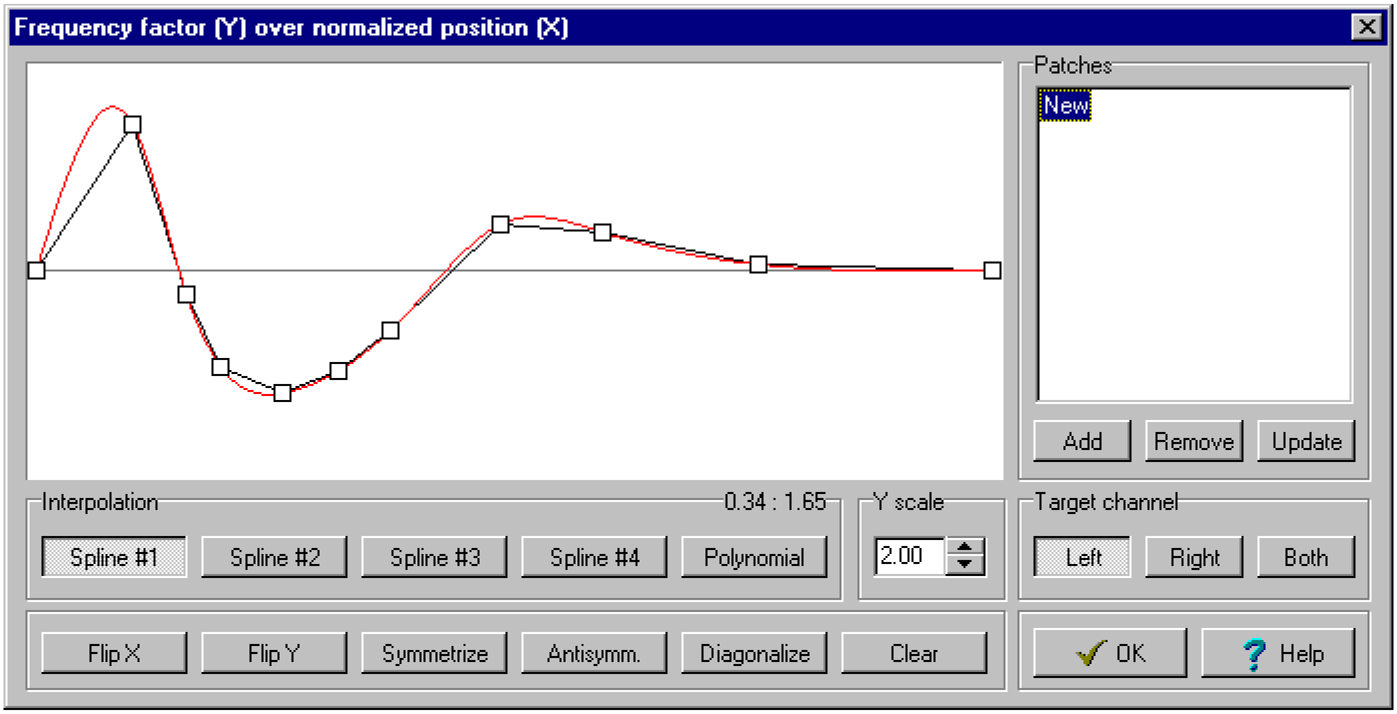

The envelope shape for the frequency factor is entered and edited in the graph. You also have the editor buttons in the bottom box at your disposal.

Use the **Y scale** box to set the overall scale for the frequency factor.

By default, envelope values between nodes are computed by linear interpolation. If a button in the **Interpolation** box is depressed, two envelope curves will be displayed: the usual, linearly interpolated one (in black) and the selected, non-linear one actually imposed on the data (in red). You can choose between five different types of non-linear interpolation.

The cursor position display in the upper right corner of the **Interpolation** box is read as "normalized position : frequency factor" (time positions are normalized, i.e. 0.00 denotes the start and 1.00 the end of the selected section).

Selecting **Left** or **Right** causes only one channel to be modulated. The resulting difference in duration is compensated by padding the shorter channel with silence, leaving the overall file length unchanged.

The **Target channel** setting is ignored for monophonic files.

You can save all settings to (and restore them from) your own modulation Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Modulate** will be executed.

Click 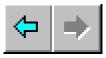 to abort and revert to the old settings.

Pitch shift

Pitch shift is used to change the frequency of the selected section while leaving its duration unchanged.

Acid WAV's pitch shifter uses an advanced, Fourier-based algorithm primarily aimed at samples of music instruments containing a dominating frequency (or set of harmonic frequencies).

Right-click **Pitch shift** to open its settings dialog:

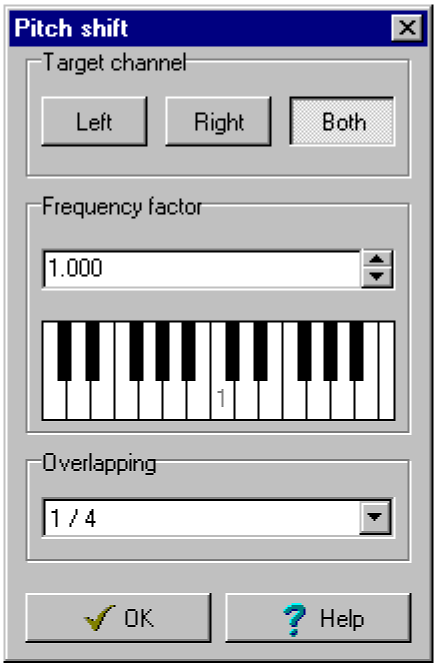

Use the **Frequency factor** box to set the amount of pitch shifting. Values between 0.5 (causing the frequency to be reduced by half) and 2.0 (causing the frequency to be doubled) are allowed. The graphic keyboard can be used to enter frequency factors corresponding to multiples of semitone shifts (useful e.g. when creating patches for samplers).

The **Overlapping** parameter determines the amount of cross-fading between successive sound chunks (like all functions, the pitch shifter reads and writes sound data in chunks a few kB in length). A setting of 1 / 1 means that the first half of each chunk is cross-faded with the second half of the previous chunk, making all chunks (except for the first and the last one) completely cross-faded with other chunks. A setting of 1 / 2 means that the first and the last first quarter of each chunk is cross-faded with other chunks, leaving the central half untouched, and so on. More overlapping tends to improve the result by smoothing out the transition between chunks. It also increases the computation time.

Selecting **Left** or **Right** causes only one channel to be pitch-shifted. The **Target channel** setting is ignored for monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Pitch shift** will be executed.

Click 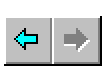 to abort and revert to the old settings.

Reduce noise

Reduce noise is used to subtract the frequency components in the clipboard from the selected section.

The most common application (and the reason for the name) is the removal of background noise (sampled into the clipboard with Copy from a silent section in the file). However, frequency subtraction can also be used to create interesting effects. In principle, any sound can be sculpted out of white noise using this method (subtractive synthesis).

Right-click **Reduce noise** to open its settings dialog:

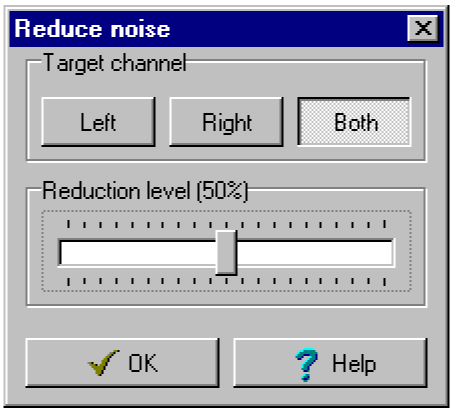

The **Reduction level** can be regarded as an amplification factor applied to the clipboard prior to subtraction. At 100%, the frequency components in the clipboard are fully subtracted from the selected section.

Selecting **Left** or **Right** causes only one channel to be affected. The **Target channel** setting is ignored for monophonic files.

Click **OK** to confirm the new settings. If the dialog was called up from the main window, **Reduce noise** will be executed.

Click 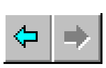 to abort and revert to the old settings.

Synthesizing new sounds

The **Synthesis** group contains the following functions:

- Draw wave
- Additive synth
- Analog synth
- FM synth
- Karplus-Strong algorithm

Draw wave

Draw wave is used to synthesize a user-defined waveform. The standard insertion rules apply.

Right-click **Draw wave** to open its settings dialog:

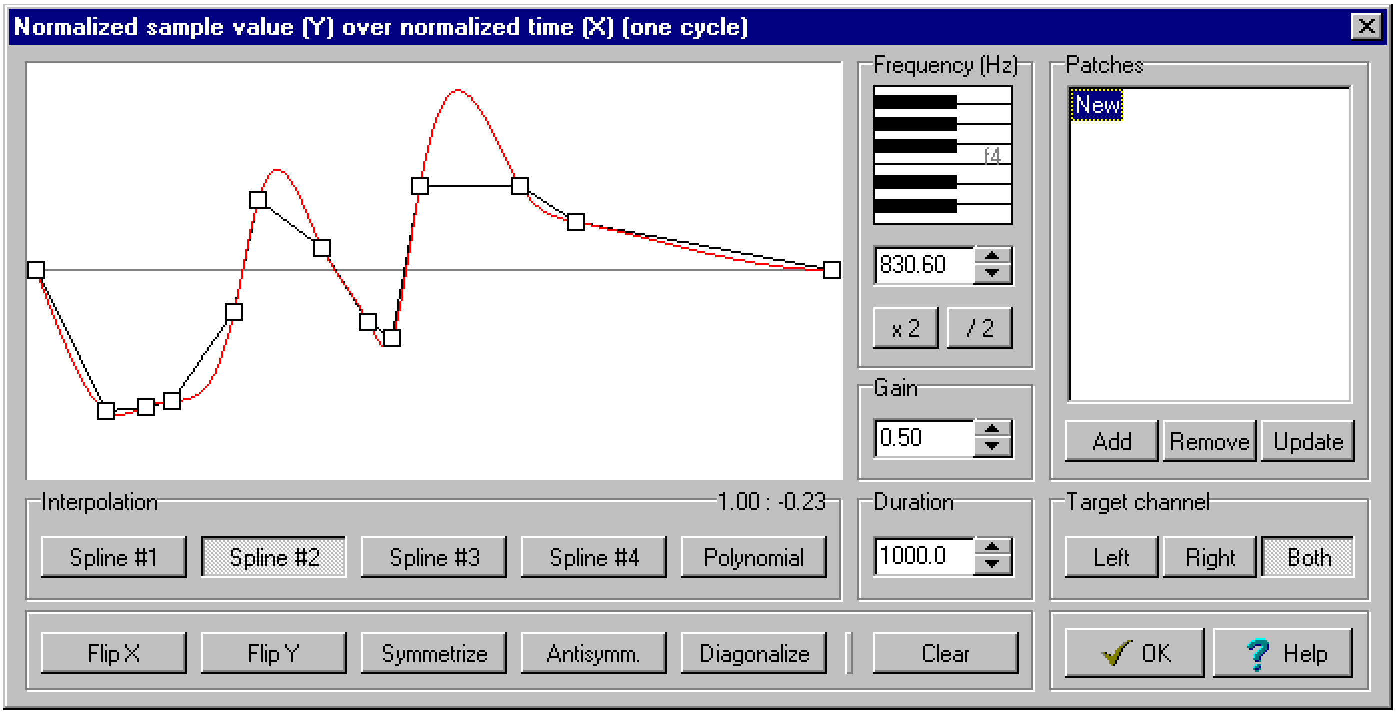

The waveform is entered and edited in the graph. You also have the editor buttons in the bottom box at your disposal.

What you draw in the graph is one cycle. Acid WAV takes care of repeating it to create a sound of the requested **Duration** and **Frequency**.

The graphic keyboard can be used to enter frequencies as notes (useful e.g. when creating patches for samplers). Use the **x 2** and **/ 2** buttons to shift the frequency value by a full octave.

When the dialog is called up from the main window, the duration is expressed in the main window's time unit.

When the dialog is called up from the script editor, the duration is expressed in milliseconds.

The cursor position display in the upper right corner of the **Interpolation** box is read as "normalized time position : normalized sample value". The time position is normalized to the duration of one cycle (i.e. the wavelength): 0.0 is the start of the cycle, 1.0 is the end. Sample values are normalized to the range -1.0 to 1.0 so as to be independent of the file format.

Use the **Gain** box to set the (normalized) amplitude scale.

By default, sample values between nodes are computed by linear interpolation. If a button in the **Interpolation** box is depressed, two waveforms will be displayed: the usual, linearly interpolated one (in black) and the selected, non-linear one actually used for synthesis (in red). You can choose between five different types of non-linear interpolation.

You can save all settings to (and restore them from) your own waveform Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Draw wave** will be executed.

Click 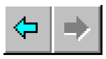 to abort and revert to the old settings.

Additive synth

Additive synth is used to mix the output from up to eight independent, analog-style oscillators. The standard insertion rules apply.

Right-click **Additive synth** to open its settings dialog:

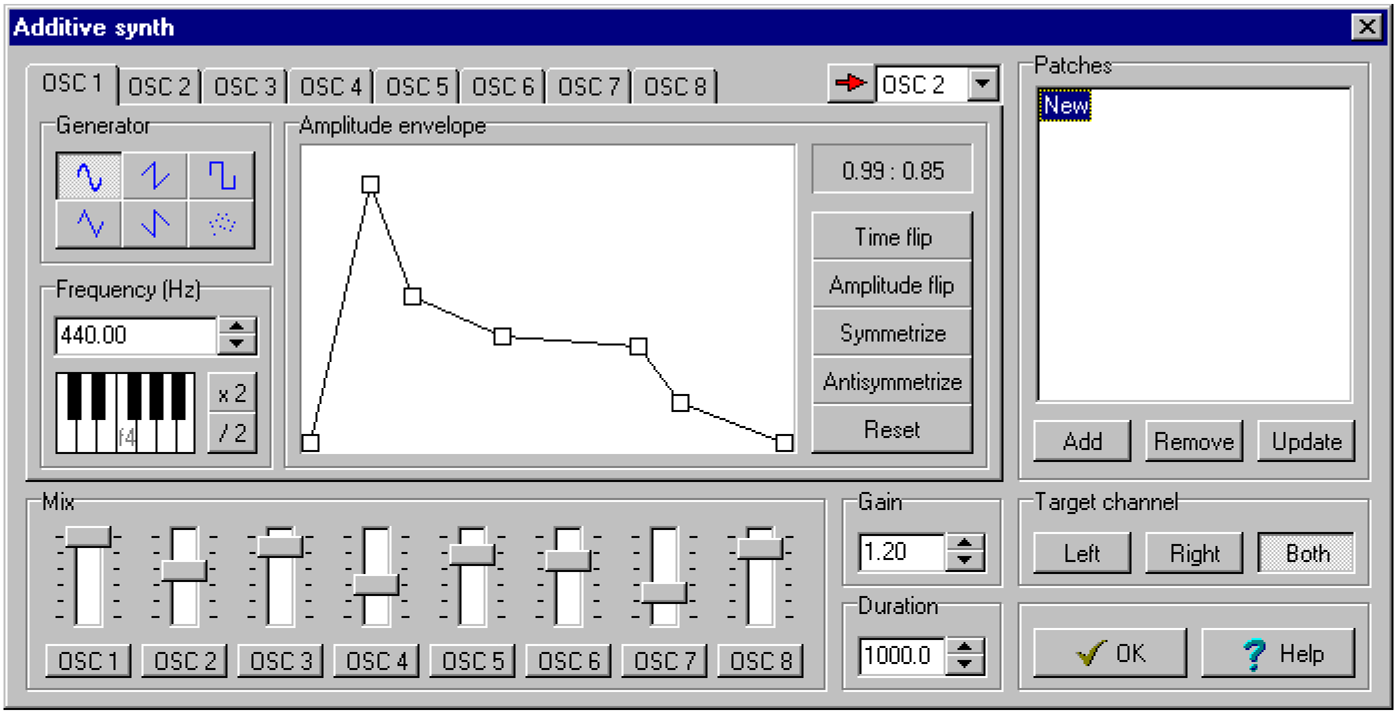

The amplitude envelope for each oscillator is entered and edited in the graph. You also have the envelope editor buttons at your disposal.

The cursor position display in the upper right corner of the **Amplitude envelope** box is read as "normalized time position : normalized amplitude". The time position is normalized: 0.0 is the start of the synthesized sound, 1.0 the end. Amplitude values are normalized to the range 0.0 to 1.0.

Use the **Generator** box to set the waveform for each oscillator:

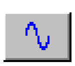 Sine wave.

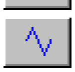 Triangle wave.

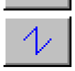 Up ramp.

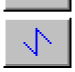 Down ramp.

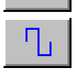 Square wave.

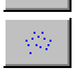 Noise.

The graphic keyboard in the **Frequency** box can be used to enter oscillator frequencies as notes (useful e.g. when creating patches for samplers). Use the **x 2** and **/ 2** buttons to shift the frequency value by a full octave.

The settings for the selected oscillator can be copied to any other oscillator using the arrow button: 

(select the target oscillator in the drop-down list).

The levels for all oscillators are set in the **Mix** box. The **OSC** button below each gauge maximizes the level for that oscillator.

You can also set an overall amplification factor in the **Gain** box.

When the dialog is called up from the main window, the sound's **Duration** is expressed in the main window's time unit.

When the dialog is called up from the script editor, the sound's **Duration** is expressed in milliseconds.

You can save all settings to (and restore them from) your own additive synth Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Additive synth** will be executed.

Click 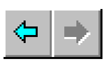 to abort and revert to the old settings.

Time flip Reverses (time-inverts) the envelope.

Amplitude flip Inverts the envelope.

Symmetrize Makes the envelope time-symmetric (a time-symmetric envelope is unaffected by reversal).

Antisymmetrize Makes the envelope time-antisymmetric (reversing a time-antisymmetric envelope is equivalent to inverting it).

Reset Deletes all but the leftmost and rightmost nodes and arranges those diagonally.

Analog synth

Analog synth is used for the creation of "vintage" synthesizer sounds. The standard insertion rules apply.

Right-click **Analog synth** to open its settings dialog:

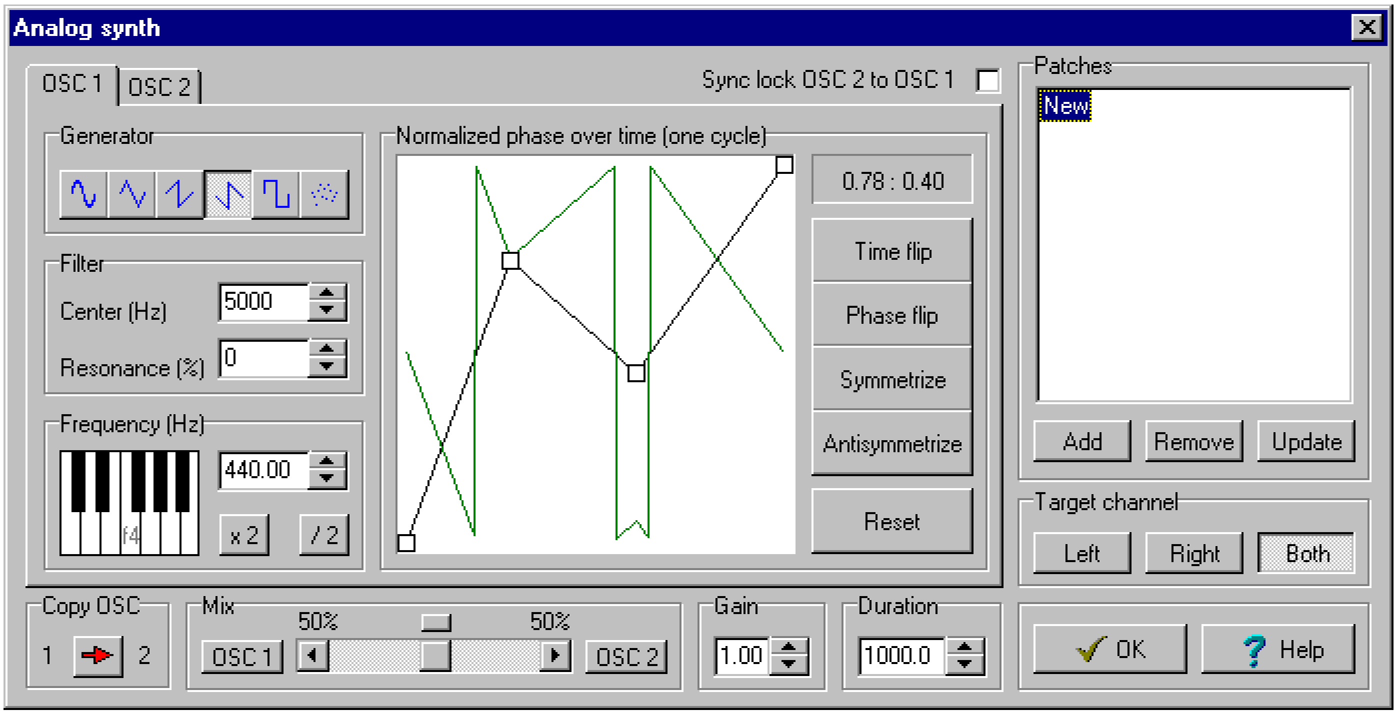

Use the **Generator** box to set the (unmodulated) waveform for the selected oscillator:

Sine wave.

Triangle wave.

Up ramp.

Down ramp.

Square wave.

Noise.

One cycle of the waveform is displayed in the graph (in green). You can modify it by editing the phase curve also displayed in the graph (in black), either directly or using the editor buttons. (Editing is disabled for noise, since noise will remain noise no matter what you do to its phase curve!)

The cursor position display is read as "normalized time position : normalized phase". Both quantities are normalized to the range 0.0 - 1.0.

Use the **Filter** box to cascade a resonant lowpass filter after the oscillator. The **Resonance** setting is the ratio between the gain at the **Cutoff** frequency and the gain at 0 Hz (1.0), expressed as a percentage. You can visualize the frequency response with the Resonate dialog. Note that when the **Resonance** is 0%, the filter reduces to plain lowpass.

The graphic keyboard in the **Frequency** box can be used to enter oscillator frequencies as notes (useful e.g. when creating patches for samplers). Use the **x 2** and **/ 2** buttons to shift the frequency value by a full octave.

The settings for the selected oscillator can be copied to the other oscillator using the arrow button in the **Copy OSC** box.

When the box labeled **Sync lock OSC 2 to OSC 1** is checked, oscillator #2 will be restarted every time oscillator #1 begins a new cycle. When the two oscillators have different frequencies, the resulting abrupt clipping of oscillator #2 can heighten the sound's brightness and introduce other interesting effects.

Use the **Mix** box to set the relative levels of the two oscillators and the **Gain** box to set the overall amplification factor.

When the dialog is called up from the main window, the sound's **Duration** is expressed in the main window's time unit.

When the dialog is called up from the script editor, the sound's **Duration** is expressed in milliseconds.

You can save all settings to (and restore them from) your own analog synth Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Analog synth** will be executed.

Click 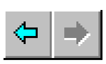 to abort and revert to the old settings.

Imagine a piece of magnetic tape containing a recording of exactly one cycle of a periodic signal. Playing the tape in a loop will recreate the sound. The position of the tape recorder's head over the tape is the phase. The phase curve is a plot of this position over the time required to loop through the tape once (i.e. to play one cycle of the sound).

If the phase curve is a straight diagonal line, the tape is running at constant speed. If the slope varies, different parts of the tape are being played at different speeds, and the resulting waveform is undergoing **phase distortion**. If the slope is positive, the tape is running forward. If the slope is negative, the tape is running backward (the sound is being reversed).

In audio engineering, phase distortion is usually a problem. It was first turned into a feature in the mid-80s, when the engineers at Casio used it in the CZ synthesizer series to create analog-like waveforms from a sine lookup table. You can use Acid WAV to emulate the old CZ sound by selecting sine waves as your **Generator** waveforms.

Time flip Reverses (time-inverts) the phase curve

Phase flip Inverts the phase curve.

Symmetrize Makes the phase curve time-symmetric (a time-symmetric phase curve is unaffected by reversal).

Antisymmetrize Makes the phase curve time-antisymmetric (reversing a time-antisymmetric phase curve is equivalent to inverting it).

Reset Deletes all but the leftmost and rightmost nodes and arranges those diagonally.

FM synth

FM synth is used to create sound by frequency modulation. The standard insertion rules apply.

Right-click **FM synth** to open its settings dialog:

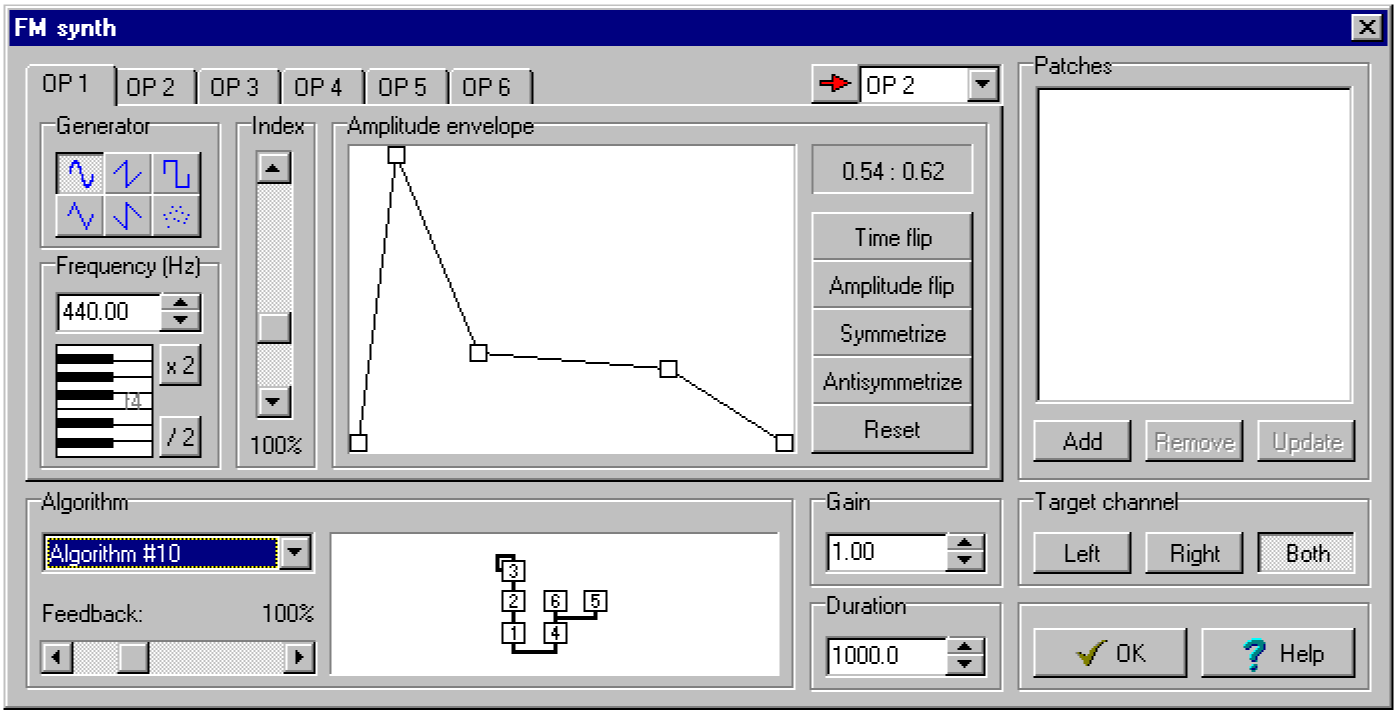

The amplitude envelope for each operator is entered and edited in the graph. You also have the envelope editor buttons at your disposal.

The cursor position display in the upper right corner of the **Amplitude envelope** box is read as "normalized time position : normalized amplitude". The time position is normalized: 0.0 is the start of the synthesized sound, 1.0 the end. Amplitude values are normalized to the range 0.0 to 1.0.

Use the **Generator** box to set the (unmodulated) waveform for each operator:

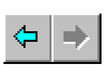 Sine wave. This is the only waveform used in most FM hardware, from cheap sound card chips for PCs to Yamaha's classic DX series synthesizers.

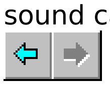 Triangle wave.

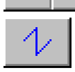 Up ramp.

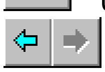 Down ramp.

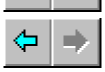 Square wave.

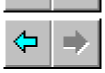 Noise.

The graphic keyboard in the **Frequency** box can be used to enter operator frequencies as notes (useful e.g. when creating patches for samplers). Use the **x 2** and **/ 2** buttons to shift the frequency by a full octave.

The **Index** setting determines the modulation depth for the operator. At 0%, the operator is unmodulated. At 100%, the modulator input is used undiminished. At higher settings, it's amplified.

The settings for the selected operator can be copied to any other operator using the arrow button:  (select the target operator in the drop-down list).

The drop-down list in the **Algorithm** box lets you choose between 32 FM signal and modulator routes. The algorithms are the same as in Yamaha's DX-7 synthesizer. All operator diagrams are read from top to bottom: in the picture above, operator #3 is modulating itself (with a **Feedback** setting of 100%) as well as operator #2. Operator #2 is in turn modulating operator #1. Operators #6 and #5 are unmodulated. Their output is mixed and used to modulate operator #4. The final output of the algorithm is obtained by mixing the outputs from operators #1 and #4.

You can set an overall amplification factor for the FM synth in the **Gain** box.

When the dialog is called up from the main window, the sound's **Duration** is expressed in the main window's time unit.

When the dialog is called up from the script editor, the sound's **Duration** is expressed in milliseconds.

You can save all settings to (and restore them from) your own FM synth Patches.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **FM synth** will be executed.

Click 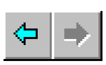 to abort and revert to the old settings.

Karplus-Strong algorithm

Karplus-Strong is used to synthesize the sound of a plucked string. The standard insertion rules apply.

The principal attraction of this synthesis method is its partial randomness. Two Karplus-Strong sections created with identical settings are unlikely to ever sound exactly the same.

Right-click **Karplus-Strong** to open its settings dialog:

The graphic keyboard in the **Frequency** box can be used to enter frequencies as notes (useful e.g. when creating patches for samplers). Use the **x 2** and **/ 2** buttons to shift the frequency value by a full octave.

The **Start** and **End** gain settings are used to control the overall amplitude and decay rate of the sound. A setting of 1.0 yields full amplitude.

The **Smoothing level** determines how fast high frequency components are attenuated relative to low frequency components. A higher setting results in faster high frequency attenuation.

When the dialog is called up from the main window, the sound's **Duration** is expressed in the main window's time unit.

When the dialog is called up from the script editor, the sound's **Duration** is expressed in milliseconds.

Click **OK** to accept the new settings. If the dialog was called up from the main window, **Karplus-Strong** will be executed.

Click 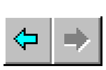 to abort and revert to the old settings.

Setting program options

Apart from the settings for individual editing and synthesis functions, Acid WAV lets you customize the following items:

- Max Open / Reset frame
- Open settings
- Output format
- Record settings
- Temporary storage
- Horizontal zoom

Max Open / Reset frame

Right-clicking the **Reset** button brings up the **Max Open / Reset frame** window:

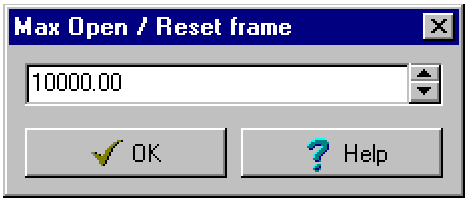

Use this window to specify the default frame length for the wave display in terms of the currently selected time unit (set in the time unit box). The default frame length is the max size of the file section shown in the wave display when you open a file or click Reset.

Click **OK** to use the value displayed in the edit box, to cancel.

The default frame length can also be set in the Open Settings window.

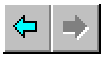

Open settings

Right-clicking the 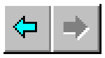 button in the toolbar brings up the **Open settings** window:

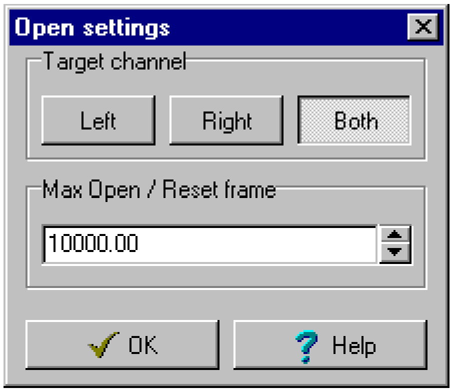

Use the **Target channel** box to specify which channel(s) you want to be loaded (if not both, the other channel is silenced). This setting is ignored for monophonic files.

Use the **Max Open / Reset frame** to specify the default frame length for the wave display in terms of the currently selected time unit (set in the time unit box). The default frame length is the max size of the file section shown in the wave display when you open a file or click Reset.

The default frame length can also be set in the Max Open / Reset frame window.

Click **OK** to make the new settings take effect and load a file, 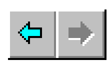 to cancel.

Output format

Right-clicking the button in the toolbar brings up the **Output format** window:

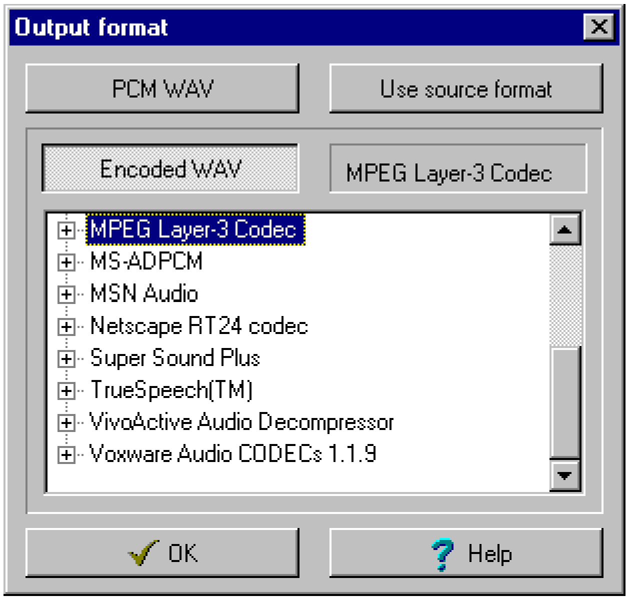

Click **PCM WAV** to select plain Pulse Code Modulation (i.e. no encoding). This is the basic format supported by all WAV readers and players.

Click **Use source format** to revert to the original format of the current file. If the file was synthesized or recorded rather than loaded, the original format is plain PCM.

Click an entry in the ACM codec (coder / decoder) list to select an encoded (usually compressed) WAV format.

Click **OK** to make the new settings take effect and save the current file, to cancel.

Acid WAV uses the Windows Audio Compression Manager to encode and decode WAV files. For a format to be available, an ACM codec supporting that format must be installed on your system. Several codecs come with Windows. Others are installed by third party applications and add-ons such as Netscape Navigator and NetShow (e.g. Microsoft's MPEG Layer-3 codec).

If an ACM codec supports several formats, it is preceded by a ⊞ icon in the list displayed by the **Output format** window. Clicking this icon causes all supported formats to be displayed. You can force Acid WAV to use a specific format, but it's usually best to just select a codec and let the Audio Compression Manager pick the most suitable format based on the word size, number of channels and sampling rate of the file.

ACM codecs can usually only handle certain combinations of word size, number of channels and sampling rate. If the codec you have selected can't handle your file, try modifying the file in a way fitting one of the codec's formats, e.g. by converting from 16 to 8 bits or from stereo to mono. See Changing file properties for details.

Record settings

Right-clicking the 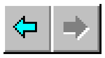 button in the toolbar brings up the **Record settings** window:

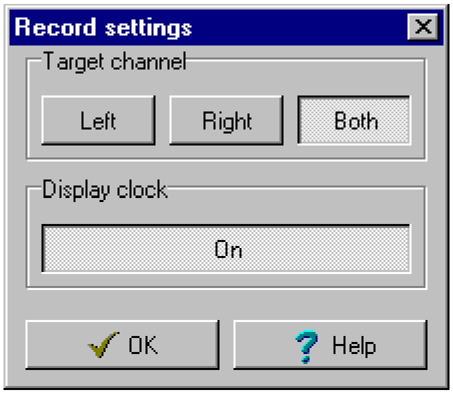

If the **Left** or **Right** buttons are depressed, recording will only be performed on the selected channel (be careful!). The **Target channel** setting is ignored when recording monophonic files.

When **Display clock** is **On**, a time and filesize box is inset in the wave display while recording:

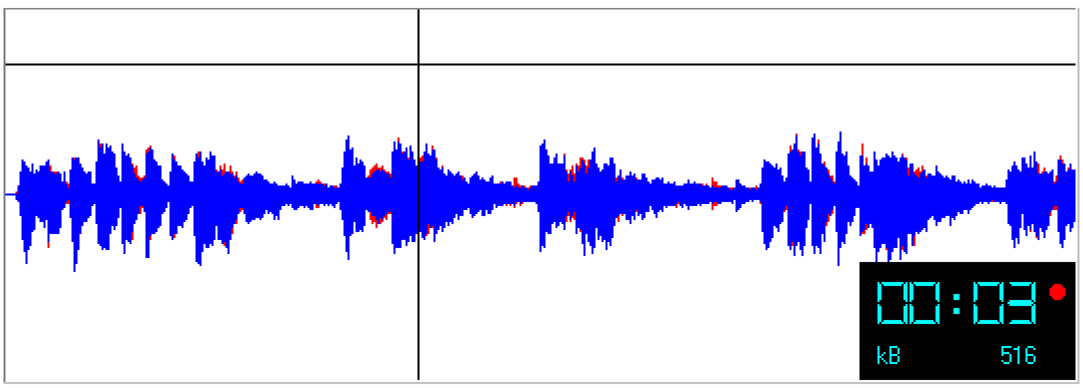

Click **OK** to make the new settings take effect and start recording, 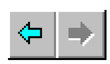 to cancel.

Temporary storage

Right-clicking the history group 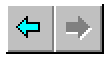 in the toolbar (even if it's grayed out) brings up the **Temporary storage** window:

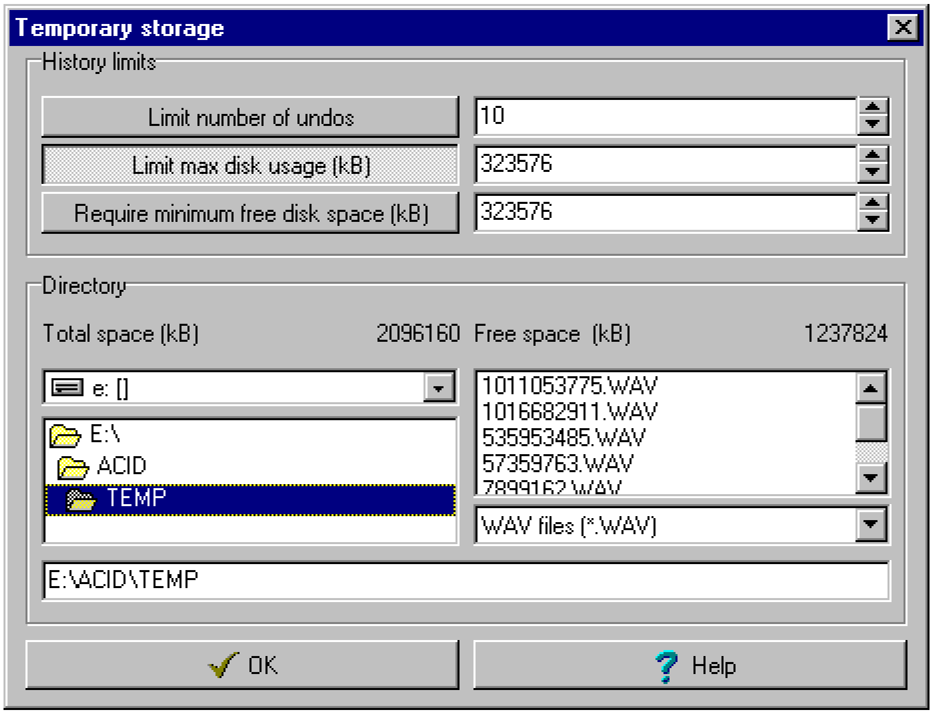

In this context, "temporary storage" essentially means "editor file". Keeping old editor files is what allows Acid WAV to undo and quickly redo operations when you click the buttons in the history group.

Since editor files can get quite large, you will probably want to limit Acid WAV's historical memory. This is done in the **History limits** box: you can put a limit on the number of undos (i.e. of old edit buffers to keep) or on disk space (either space available for temporary storage or "headroom", i.e. the minimally acceptable amount of free disk space). You can also combine all three kinds of limits.

In the picture above, Acid WAV is limiting its disk usage for temporary storage to 323576 kilobytes.

The **Directory** box lets you set the disk and directory to use for temporary storage. By default, Acid WAV uses the TEMP subdirectory in its own home directory.

Right-clicking the **Horizontal zoom** button in the selection box (even if it's grayed out) causes the following one-item menu to pop up:

Click the item to check / uncheck it. When it's checked, clicking the **Horizontal zoom** button causes Acid WAV to unselect the marked file section after fitting the frame length to it.

Automating tasks with scripts

Right-clicking 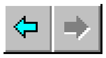 in the toolbar brings up Acid WAV's scripts window:

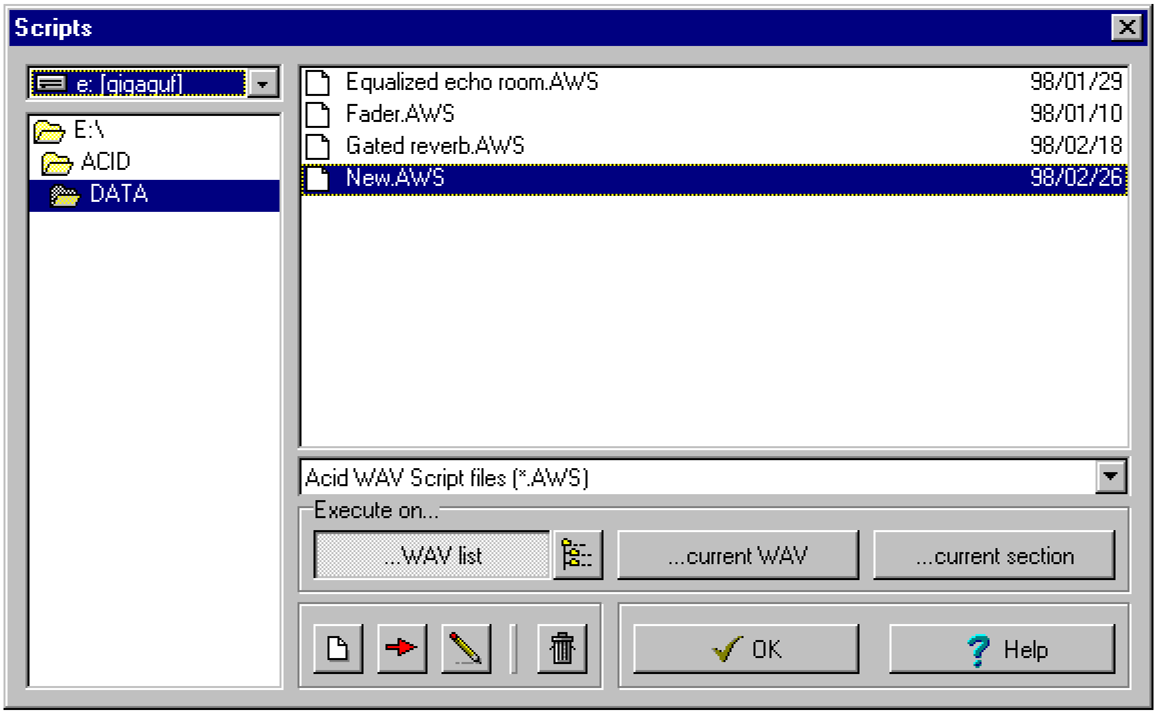

Use this window to

- create, copy and delete scripts,
- edit scripts,
- set the default script,
- execute scripts.

Closing the **Scripts** window with **OK** causes the selected script (if any) to be executed. Click 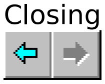 to exit without executing any script.

A script is a file which

- describes a sequence of Acid WAV operations, complete with all relevant settings, and

- can be executed as a single command.

Using scripts, you can

- speed up interactive work with your own Acid WAV "functions" (perform frequently recurring and lengthy tasks with just one click) and

- let Acid WAV run the same operation sequence on multiple files without your intervention (batch processing).

Script files are denoted by the extension AWS (for Acid WAV Script). Although you can keep them wherever you want, the recommended location is Acid WAV's DATA subdirectory.

Click

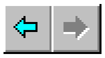 to create a new script.

When a script is selected, you can also click

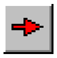 to copy it,
 to delete it and
 to edit it.

Editing scripts

Clicking  in the **Scripts** window causes the selected script to be opened in Acid WAV's script editor:

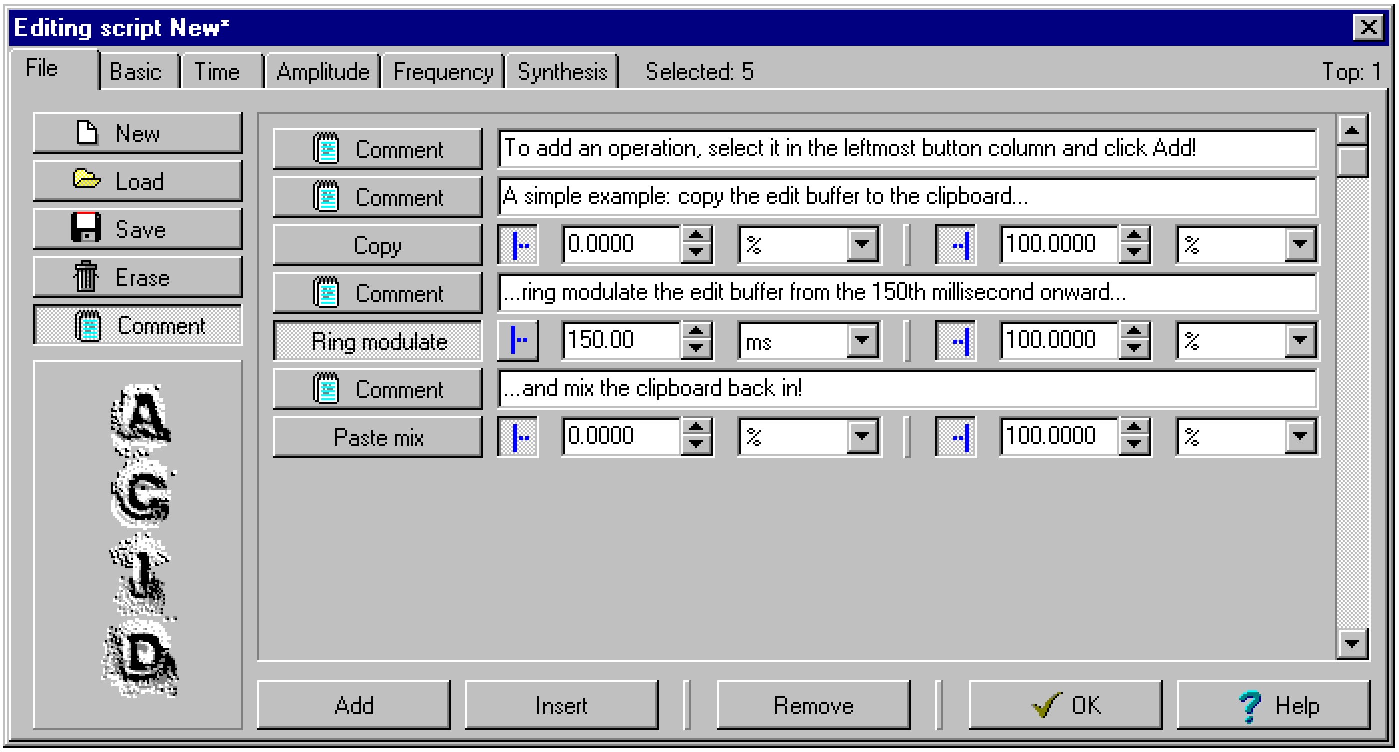

In this picture, the title bar tells us that we are editing a script called "New". The familiar asterisk trailing the script's name means that the script has been modified. Nothing is written to disk until **OK** is clicked. Closing the script editor with 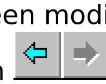 causes all changes to be discarded.

Do not be confused by the buttons under the **File** *tab! The leftmost button column is the "palette" where you select operations to* **Add** *or* **Insert** *into the script. It is not used to load or save scripts.*

Also note that the source files used in batch runs are loaded and saved automatically by Acid WAV. There is no need to put **Load** and **Save** operations in scripts for that. The primary purpose of **Save, Load** (and **Erase**) in scripts is to supplements the clipboard, i.e. for temporay storage of intermediate results.

Anatomy of a script

The central part of the editor contains the script itself.

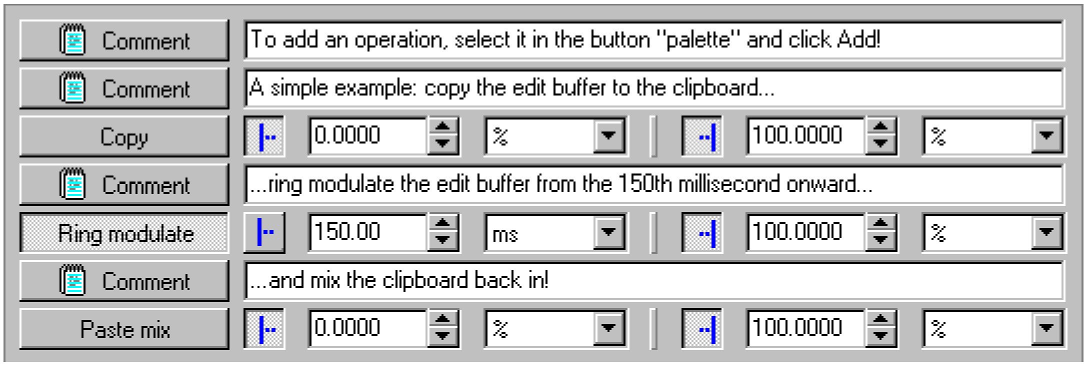

Each script line consists of

- a button labeled with the name of an operation and
- operation-dependent data fields (text for comment lines, start and end positions for most editing and synthesis functions, file names for **Load**, **Save** and **Erase**).

The first two lines in the picture are comments. The third line contains the first actual operation (**Copy**).

The operation button in each script line fills two purposes:

- When it's depressed, it tells the editor where to **Add** and **Insert** new lines, as well as which line to **Remove** (all done with the buttons at the bottom of the editor window):

Only one operation button can be depressed (i.e. only one script line can be selected) at any one time. The number of the selected line (if any) is displayed below the title bar, along with the number of the topmost visible line (below the 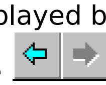 button).

If no line is selected, new lines are **Add**ed after the last one and **Insert**ed before the first one; **Remove** is grayed out.

- Right-clicking the operation button brings up the line's settings window (if one exists). Each script line has its own settings. Editing those settings does not affect the settings of other script lines or of the main Acid WAV window.

Most settings windows you will see in the script editor are identical to those you are used to from Acid WAV's main window. The only (minor) exceptions are Load, Save and Insert file.

Start and end positions

The script editor inserts start and end position controls for all operations affected by file section selections:

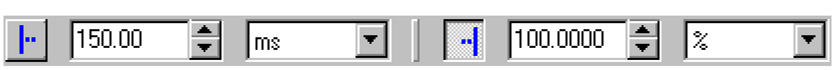

When a script is executed on a whole file (**...current WAV** or **...WAV list** in the **Scripts** window) the following rules apply:

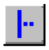 Sets the start position to the first sample in the file.

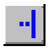 Sets the end position to the last sample in the file.

Positions expressed in samples or milliseconds are relative to the first sample in the file.

Positions expressed as percentages are converted to absolute values by multiplication with the file length / 100.

When a script is executed on a section (**...current section** in the **Scripts** window, 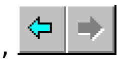 in the toolbar) Acid WAV tries to recompute all positions relative to that section.

Load

Unlike the ordinary Open settings window, the script editor's **Load** settings window does not have a Max Open / Reset frame box. At runtime, scripts use the main window's max frame setting.

The source files used in batch runs are loaded automatically by Acid WAV. There is no need to put **Load** *operations in scripts for that purpose.*

The primary purpose of **Save, Load** and **Erase** in scripts is to supplement the clipboard, i.e. for temporay storage of intermediate results.

Insert file

Apart from the controls also contained in the main window's Insert file dialog, the script editor's counterpart also requires that you specify the name of the file to be inserted:

The file name must be specified, but you may leave out the directory. In that case, Acid WAV will look for the file in whatever directory happens to be the default one when the script is executed.

Output format

Unlike the Output format window brought up by right-clicking 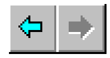 in Acid WAV's toolbar, the **Output format** window displayed by the script editor and by the **WAV list** window does not have a **Use source format** button.

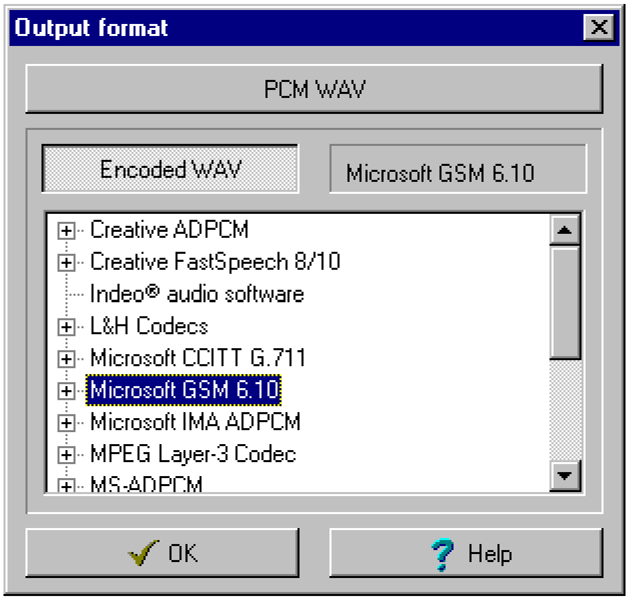

If you want to **Save** from within a script, you must specify the destination format explicitly.

The results of batch runs are saved automatically by Acid WAV.There is no need to put ***Save*** *operations in scripts for that purpose.*

The primary purpose of **Save, Load** and **Erase** in scripts is to supplement the clipboard, i.e. for temporay storage of intermediate results.

Before the first script operation is executed,
the following values are saved:

START = source section's start position
LENGTH = source section's length
DISTANCE = distance from source section's end position to last sample in file

Then...

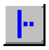
Sets the start position to START.


Sets the end position to (last sample in file - DISTANCE).

Positions expressed in samples or milliseconds are relative to START.

Start positions expressed as percentages are converted to absolute values using the formula START + (percentage * LENGTH / 100).

End positions expressed as percentages are converted to absolute values using the formula (last sample in file - DISTANCE) - (LENGTH * (percentage - 100) / 100).

Selecting a script in the **Scripts** window makes it the default script, i.e. the script to be executed on the selected file section when you click 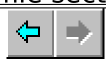 in the main window's toolbar.

Executing scripts

The fastest way to execute a script is to click 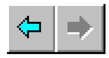 in the main window's toolbar. This causes the default script (the currently selected script in the **Scripts** window) to be executed on the selected file section.

Executing a script on the selected file section means that all start and end positions contained in the script are interpreted relative to the file section rather than to the entire file. For instance, if the script calls for a function to be applied from a start position of 10 ms, this will be interpreted as "from a start position 10 ms into the selected file section".

If there is no default script, clicking 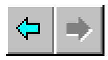 causes the **Scripts** window to open.

The **Scripts** window lets you execute scripts in three different ways:

- To execute the script on the selected file section, click **...current section** in the **Execute on...** box, followed by **OK**.
- To execute the script on the entire file (i.e. ignoring start and end markers) click **...current WAV** in the **Execute on...** box, followed by **OK**.
- To execute the script on several files (batch processing) click **...WAV list** in the **Execute on...** box, followed by **OK**.

For the **...WAV list** button to be enabled, you must have provided Acid WAV with a list of source files using the WAV list window.

The **WAV list** window is opened by clicking 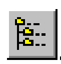.

WAV list

Clicking 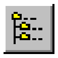 in the **Scripts** window brings up the **WAV list** window:

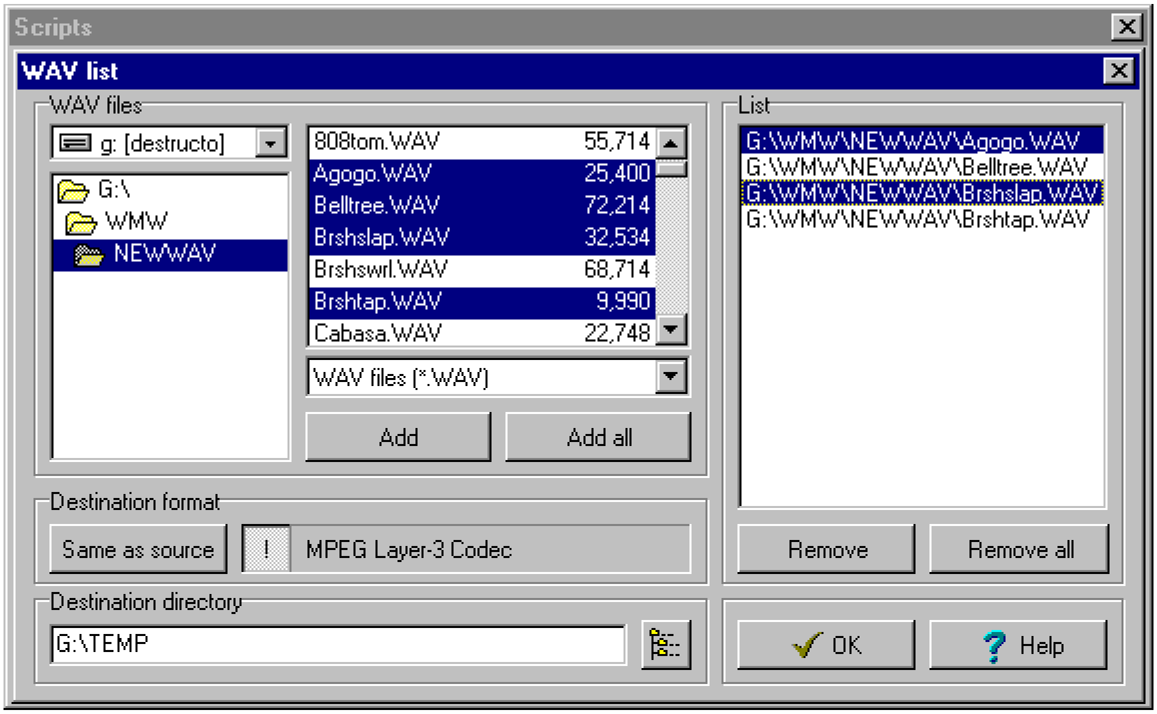

Use the controls in the **WAV files** box to select files and **Add** them to the **List**. Use the controls in the **List** box to **Remove** files from the list. Multiple selections are supported (press the Shift key while clicking to select a range; press Ctrl while clicking to select additional individual files).

Use the **Destination format** box to set the ACM encoding (if any) to be used in the destination files. When **Same as source** is depressed, destination files are saved using the same encoding as the corresponding source files. Click the **!** button or the encoding display to open the Output format window.

Use the **Destination directory** box to tell Acid WAV where to put destination files. Click 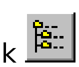 to browse the directories on your hard disk.

Leaving the **Destination directory** edit box blank causes the source directory to be used as the destination directory, too.

Destination files are given the same name as source files, so make sure to use separate source and destination directories unless you want the source files to be overwritten!

Ordering overview

The freely distributable Acid WAV Evaluation Package lets you use all Acid WAV functions, but not all at the same time. For unlimited functionality, you need to purchase a personal key from Polyhedric Software. The key comes in the form of a small binary file and is delivered either as an e-mail attachment (this is the fastest and least expensive option) or on a standard 3.5 inch diskette (recommended only if you don't have e-mail access or if your e-mail system can't handle binary attachments).

Please refer to the License Agreement for legal details regarding the use and distribution of Acid WAV and Acid WAV keys.

Polyhedric Software is letting Kagi (a specialized order processing firm) handle all orders for Acid WAV keys. Orders can be submitted to Kagi through an online form, by e-mail, by fax or by snail mail. The online form is the fastest route, snail mail the slowest. Kagi accepts cash (all major currencies), checks, money orders, credit cards (VISA, MasterCard, American Express, Discover or Diners Club), First Virtual and invoices. If you need a paper receipt you can have one sent to you.

An Acid WAV key is only $45!

- To submit your order online, you need a credit card. A secure form is available for browsers supporting SSL (e.g. newer versions of Netscape Navigator and Internet Explorer). Data entered on the secure form is encrypted before it's sent over the Internet.

 Click here to submit your order online.

- To submit your order by e-mail, fax or snail mail, you need to run Kagi's "Register" wizard. It will prompt you for all the necessary information and encode any credit card data for better security.

 Click here to submit your order by e-mail, fax or snail mail.

If Windows Help complains about not being able to find Register.EXE, you can launch it from Acid WAV's folder (click the icon labeled "Order").

Send any questions about ordering Acid WAV keys to sales@polyhedric.com.

Kagi

1442 - A Walnut Street #392

Berkeley, CA 94709-1405

USA

Fax: +1 510 652 6589
http://www.kagi.com
kagi@kagi.com

To order your Acid WAV key on the web, you need a credit card (VISA, MasterCard, American Express, Discover or Diners Club). A secure form is available for browsers supporting SSL (e.g. newer versions of Netscape Navigator and Internet Explorer). Data submitted through the secure form is encrypted before it's sent over the Internet.

 Click here to submit your order on the web.

To order your Acid WAV key by e-mail, fax or snail mail, you need to run Kagi's "Register" wizard. It will prompt you for all the necessary information and encode any credit card data for better security.

Click here to submit your order by e-mail, fax or snail mail.

If Windows Help complains about not being able to find Register.EXE, you can launch it from Acid WAV's folder (click the icon labeled "Order").